DOCUMENT RESUME

ED 371 721 IR 016 641

TITLE	The DOE Program in HPCC: High-Performance Computing and Communications.
INSTITUTION	Department of Energy, Washington, D.C. Office of Energy Research.; Department of Energy, Washington, DC. Office of Scientific Computing.
REPORT NO	DOE/ER-0536
PUB DATE	Mar 93
NOTE	104p.
AVAILABLE FROM	National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161.
PUB TYPE	Reports - Descriptive (141)
EDRS PRICE	MF01/PC05 Plus Postage.
DESCRIPTORS	*Computer Oriented Programs; Computer Uses in Education; Cooperation; Federal Legislation; Federal Programs; Government Role; Objectives; Program Descriptions; Program Implementation; Student Recruitment; *Systems Development; *Technological Advancement; *Telecommunications
IDENTIFIERS	*High Performance Computing; High Performance Computing Act 1991; National Research and Education Network; Supercomputers

ABSTRACT

This document reports to Congress on the progress that the Department of Energy has made in 1992 toward achieving the goals of the High Performance Computing and Communications (HPCC) program. Its second purpose is to provide a picture of the many programs administered by the Office of Scientific Computing under the auspices of the HPCC program. These include work on the development of high-performance computing systems, the Advanced Software Technology and Algorithms (ASTA) program, the National Research and Education Network (NREN), and the Basic Research and Human Resources (BRHR) program. In accordance with the HPCC Act of 1991, the Department of Energy has established two collaborative consortia at Los Alamos (New Mexico) and Oak Ridge (Tennessee) to work on identified "Grand Challenge" problems in a number of areas. Six figures are included. (SLD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

DOE/ER-0536

IR

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

ED 371 721

The DOE Program in

HPCC

High-Performance Computing and Communications

IR016641

March 1993

US DEPARTMENT OF ENERGY
Office of Energy Research
Office Of Scientific Computing

BEST COPY AVAILABLE

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831; prices available from (615) 576-8401.

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161.

DOE/ER-0536

The DOE Program in

HPCC

High-Performance Computing and Communications

March 1993

US DEPARTMENT OF ENERGY
Office of Energy Research
Office Of Scientific Computing
Washington, D.C. 20585

CONTENTS

Section 1: Introduction

Overview 2
Descriptions of Funding Codes for HPCC Categories 7
Description of HPCC Program 7
Program Management 10
Federal Register Announcement 10

Section 2: Highlights

Ames National Laboratory 16
Argonne National Laboratory 17
Fermi National Laboratory 18
Lawrence Berkeley Laboratory 20
Lawrence Livermore National Laboratory 22
Los Alamos National Laboratory 26
Oak Ridge National Laboratory 31
Sandia National Laboratories 35

Section 3: Program Summaries

High Performance Computing Systems 38
Advanced Software Technology and Algorithms 45
Advanced Computation, Communications Research, and Associated Activities 59
Basic Research and Human Resources 67
University Summaries 86

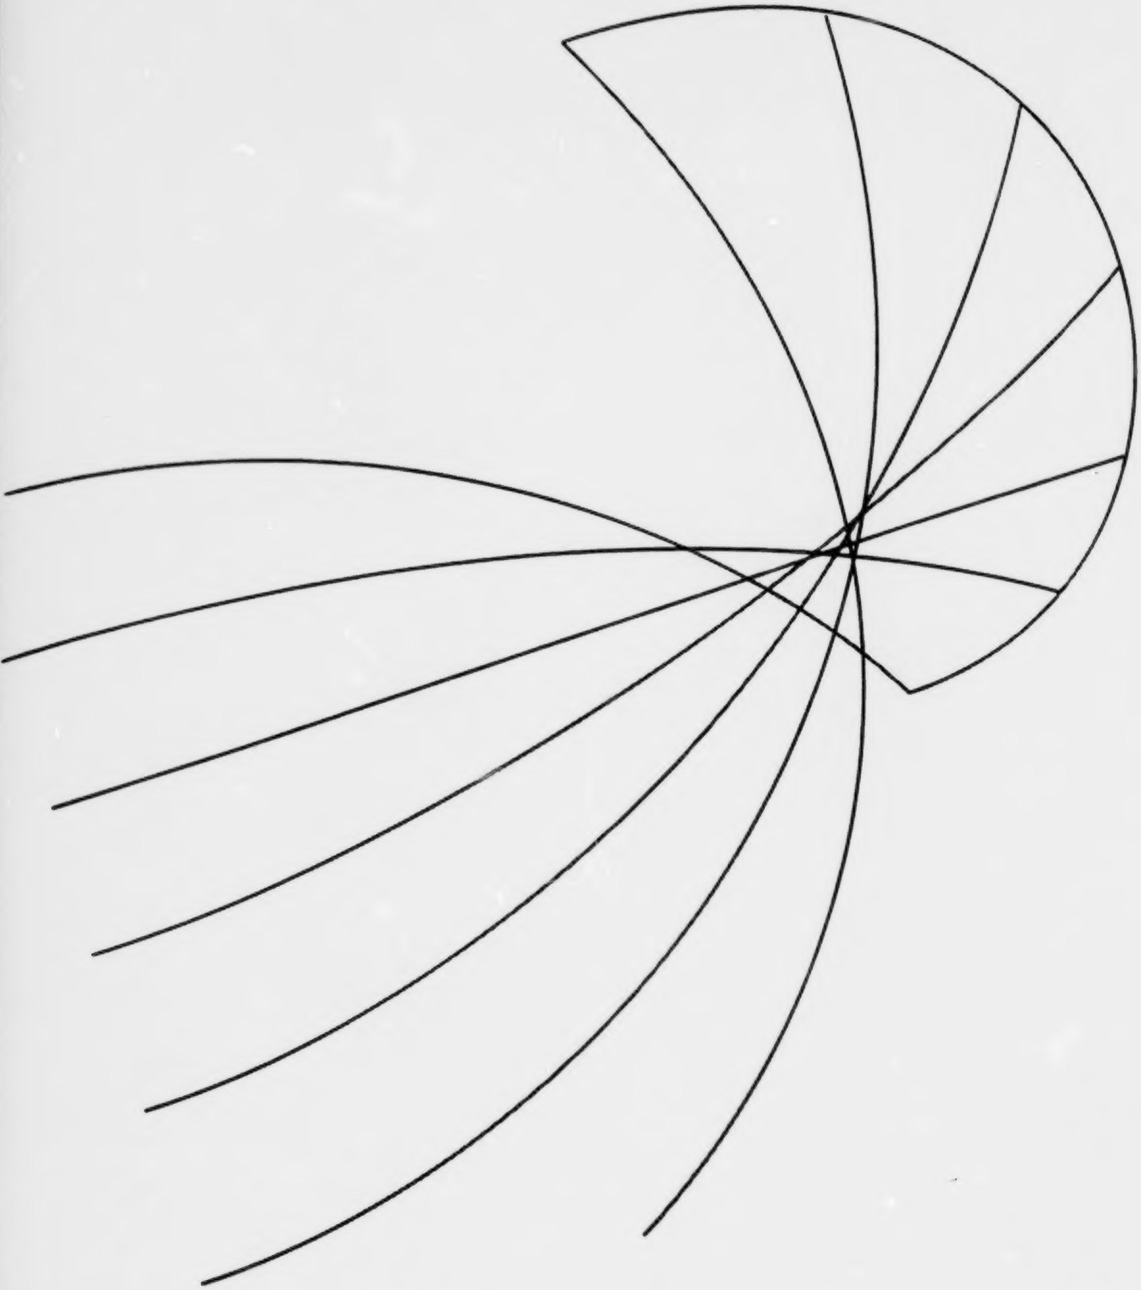

Section 1
Introduction

OVERVIEW

INTRODUCTION

This document is intended to serve two purposes. Its first purpose is to report to Congress, as mandated in the High-Performance Computing Act of 1991, the considerable progress that the Department of Energy has made in 1992 toward achieving the goals of the High-Performance Computing and Communications (HPCC) Program. Its second purpose is to provide a snapshot of and an introduction to the many programs administered by the Office of Scientific Computing under the auspices of the HPCC Program.

Background

The original Applied Mathematical Sciences (AMS) Program was started in the early fifties at the suggestion of John von Neumann to investigate the applicability of the new digital computers in solving complex scientific and technological problems of importance to the Department of Energy and the nation. Over the last forty years the AMS Program has changed considerably, but it has never lost sight of its original purpose, which is to support mathematical and computational research that facilitates the use of the latest high-performance computer systems in advancing our understanding of science and technology.

In 1983 the AMS Program was incorporated into the newly formed Scientific Computing Staff, which reported to the Director of Energy Research. At that time the areas of responsibility of the Scientific Computing Staff were divided into four parts: Analytical and Numerical Methods, Information Analysis Techniques, Advanced Computing Concepts, and Energy Sciences Advanced Computing. Research in applied mathematics and algorithm development was supported under the Analytical and Numerical Methods subprogram, and research in statistical methods, scientific visualization, and geometry was supported under Information Analysis Techniques.

Under Advanced Computing Concepts, work on the early evaluation and benchmarking of experimental computer hardware and software was undertaken, and the Energy Sciences Advanced Computing subprogram provided DOE-wide access to supercomputers for laboratory and university researchers through the operation of the Magnetic Fusion Energy Computer Center at Lawrence Livermore National Laboratory.

In the context of the present HPCC Program, it is particularly interesting to note that the Scientific Computing Staff was one of the pioneers in encouraging and supporting research in parallel computing. The AMS Program acquired the first experimental parallel machine in the Department of Energy (the Denelcor Heterogeneous Element Processor at Los Alamos National Laboratory), and it supported the design and construction of the first hypercube parallel computer (the Cosmic Cube at Caltech), along with other novel computer architectures, such as the Ultracomputer at New York University and the CEDAR cluster at the University of Illinois at Urbana. The AMS program continues to operate the National Energy Research Supercomputer Center at Livermore (the successor to the MFE Center), which provides supercomputer and networking services to the DOE research and educational communities.

In 1983 the Lax Report cited the successful DOE programs in computational science and it called for the federal government to provide supercomputer resources to all US scientists. As a result of this report and others like it, a number of federally sponsored programs in supercomputing research and access were begun, among them the DOE Supercomputer Initiative, the DARPA Strategic Computing Initiative, and the NASA Telescience Initiative. The National Science Foundation created supercomputer centers at several universities. It is not surprising, then, that these four agencies have come to play the leading role in the development and implementation of the HPCC Program.

High-Performance Computing and Communications Program

These and other federally sponsored efforts in supercomputing access and research culminated in the establishment of the President's High-Performance Computing and Communications Initiative, which was announced on February 5, 1991, by Dr. D. Alan Bromley, the presidential science advisor. Subsequently, the Congress passed the High-Performance Computing Act of 1991 (P.L. 102–194), which was signed into law by President Bush on December 9, 1991.

The HPCC Initiative seeks to increase and strengthen American competitiveness in the areas of national security, science, technology, industry, and education by accelerating the research, development, deployment, and acceptance of a new generation of computing, communication, and information technology tools. At the present time nine federal agencies are involved directly in the HPCC Initiative.

The HPCC program has four major components—HPCS, ASTA, NREN, and BRHR—which are described below.

The High-Performance Computing Systems (HPCS) component consists primarily of research, development, and prototype-evaluation activities that advance the capabilities of hardware technology to solve Grand Challenge problems like global climate modeling. (Grand Challenge problems are those problems whose solutions are considered critical to national needs.) Other activities include the design and implementation of resources such as data storage and retrieval that are required to solve Grand Challenge class problems.

The Advanced Software Technology and Algorithms (ASTA) component includes research and development of software tools, computational techniques and numerical algorithms for solving the Grand Challenge problems and representing the solutions in a convenient visual or graphical form. Included in this component are the activities of High-Performance Computing Research Centers, which are the intellectual homes of research on various Grand Challenge problems.

The National Research and Education Network (NREN) component is concerned with very-high-speed digital communications and the staged deployment of a multigigabit national network backbone. The national network will provide services such as full-motion video, rapid transfer of high-resolution images, real-time display of time-dependent graphics, and remote operation of experiments to institutions across the United States, acting as a uniform network interface to domestic users and a standard interface to international research and education networks.

The Basic Research and Human Resources (BRHR) component includes basic research in applied mathematics and computer science. In addition, this component supports activities in K–12, undergraduate, graduate, and postgraduate education; training; curriculum development; and enhancement of the educational and personnel infrastructure.

DOE COMPONENT OF THE FEDERAL HPCC PROGRAM

Commensurate with its new HPCC responsibilities, the Scientific Computing Staff was elevated in 1992 to program office status within the Office of Energy Research and is now the Office of Scientific Computing. The mission of the Office of Scientific Computing is to improve the ability of DOE to solve scientific and engineering problems that are critical to its mission and to the national interest by supporting research in and applications of advanced mathematical, computer, and computational sciences.

The Office of Scientific Computing has reorganized its program areas to better carry out its expanded mission under the HPCC Program. The areas of responsibility include the components of the focused HPCC Program described above, as well as the important elements of the base AMS Program that predate the HPCC Initiative.

An important aspect of this reorganization has been the Office's outreach beyond its traditional university and laboratory communities toward direct cooperation with industrial organizations and consortia. This outreach is reflected in several of the activities described below.

The increasing emphasis that the Office of Scientific Computing is placing on industrial cooperation and technology transfer is illustrated by the rapid increase in the number of CRADAs (Cooperative Research and Development Agreements) established between teams of laboratory researchers and industrial companies that stand to benefit from advances in high performance computing and communications. The twelve CRADAs sponsored by the Office in 1992 are excellent vehicles for providing companies of all sizes access to the mathematical and computational expertise of the laboratory research community, which in turn benefits from exposure to difficult and interesting problems at the cutting edge of science and technology. Significant industrial involvement was of prime importance also in the selection and funding of the Grand Challenge projects initiated in 1992 by the Office of Scientific Computing. The particular company or group of companies was considered an integral part of the proposed research effort and an equal partner in the project along with the teams of university and laboratory researchers. A final illustration of the increased industrial participation in AMS-supported activities is afforded by the National Storage Laboratory (NSL), a collaborative project created to remove storage system bottlenecks, provide much-needed storage system functionality, and facilitate the establishment of national storage system standards. Located at the Lawrence Livermore National Laboratory (LLNL), the NSL is guided by the premise that no single company, government laboratory, or research organization has the capability to confront all of the system-level issues involved in making significant advances in storage systems technology. Ten US computer companies, along with the four National Science Foundation Supercomputer Centers, are working with LLNL and other DOE laboratories to test the NSL system architecture on very-large-scale applications and provide commercial vendors with new ways of utilizing storage hardware and software more effectively.

High-Performance Computing Systems (HPCS)

The Department of Energy has an active and ongoing interest in the development of the foundations for future generations of computing hardware and software that will increase speed and capacity to the level of trillions of operations per second (teraflops). The DOE also has an interest in encouraging the development of future generations of machines in such a way that architecture advances can be integrated into effective, flexible computing environments. An important long-term goal that guides DOE research in this area is the development of high-performance computational facilities that are architecturally balanced in a way that makes them available for general DOE applications. The Department of Energy is especially interested in research that leads to increasing understanding of the architectural limitations that will shape the design of future machines and the impact on basic parallel software of architecture scale-ups that lead to higher performance.

Advanced Software Technology and Algorithms (ASTA)

Robust and efficient software and algorithms are the keys to making HPCC systems useful and economically viable. There is clear and overriding evidence from experiences with earlier computer systems that improvements in software and algorithms lead to greater gains in performance than do improvements in machine architecture. However, the computer systems of the past had relatively simple architectures, making it possible to develop software and applications algorithms without having to pay too much attention to the particular machine architecture involved. Today the utility of such simple architectures for software development has been thoroughly exploited. The complexity of current and contemplated architectures requires the extensive development of new software technologies, computational methods, and programming techniques to fully realize the enormous computing power of parallel and distributed computer systems.

Addressing this particular challenge requires a rethinking and a re-evaluation of the disciplinary and organizational relationships that delineate the traditional method of conducting research at most universities and federal laboratories. In place of an individual group of specialists in one area of science who work by themselves, there must be a team composed of specialists from the discipline science who work closely with computational scientists and computer scientists to develop and implement efficient and effective software and algorithms. This new research paradigm guided the formation in 1992 of research groups composed of laboratory and university specialists in the relevant discipline science and the computational sciences, working in some instances with scientists from industry, to attack Grand Challenge problems in areas such as global climate modeling, high-energy physics, fusion energy, materials science, structural biology, and rational drug design. A particular scientific problem drives the development of improved software and algorithms, and in turn, that area of science is enriched by the discovery of new phenomena made possible by advances in computational power and visualization.

National Research and Education Network (NREN)

In 1985 Congress mandated a study of the computer networking requirements of the research community in the United States. The preparation of such a far-reaching study, together with a recognition of the widespread use of remote access to supercomputers through federal supercomputer programs, focused the attention of the US research community and the federal agencies on the need for reliable, high-speed computer networks.

Within the Department of Energy, computer networks had been used extensively since the mid 1970s for various scientific and technological applications, primarily in the areas of fusion energy and high-energy physics. However, these networks had limited capacity and were incompatible with each other. Owing to a significant increase in networking requirements, the Energy Research (ER) community endorsed a proposal to create the Energy Sciences Network (ESNet). The ESNet has been developed to be compatible with existing networks, while incorporating a nondisruptive transition path to emerging international network standards. The ESNet is the vehicle through which the ER community has become a full partner in the Internet community of computer networks and through which the ER community will become an integral part of the National Research and Education Network.

Basic Research and Human Resources (BRHR)

The Office of Scientific Computing supports basic research in applied mathematics and computer science that underpins most, if not all, of its sponsored activities in high-performance computing and communications. While the first three components of the HPCC Program outlined above are concerned with research and implementation that have a near- or intermediate-term impact, the activities sponsored under the basic research component are necessary for the long-term health and vitality of computational science.

The base program in applied mathematics supports a broad range of activities at universities and the DOE laboratories in the complementary areas of modeling, analysis, and numerical simulation of physical phenomena that arise in energy systems. Most of the projects involve applications-driven studies of the mathematical and numerical tools required for understanding the behavior of complex discrete and continuous systems.

The base program in computer science focuses on understanding how parallel and distributed computer systems can be applied effectively to large-scale scientific problems. Supported projects include research in programming development models and tools, management and visualization of scientific data, improved libraries for parallel computers, software performance analysis techniques, and message-passing utilities to facilitate distributed computing.

Development and enrichment of the nation's human resources in applied mathematics, computer science, and computational science through investment in education and training at all levels from elementary school through college and beyond are the goals of the human resources component of the HPCC Program. The Office of Scientific Computing, in partnership with the Office of University and Science Education Programs, has in place a number of educational programs in scientific computing aimed at all grade levels, many of which are conducted in collaboration with the DOE laboratories.

Highlights

In accordance with the High-Performance Computing Act of 1991, the Department of Energy has established two Collaborative Consortia, which are in fact the High-Performance Computing Research Centers at Los Alamos and Oak Ridge National Laboratories. Each center is conducting research aimed at the solution of Grand Challenge problems whose size and complexity require high-performance computing and communication resources and the combined expertise of teams of interdisciplinary researchers. Each center has acquired a new massively parallel computer as the centerpiece of its Grand Challenge research. The Los Alamos HPCRC has acquired a Thinking Machines Corporation CM-5, and the Oak Ridge HPCRC has acquired an Intel Corporation Paragon. Both centers provide sophisticated testbed computing environments in which participants from other federal laboratories, industry, and academia can examine new hardware and software designs and explore new areas of science and technology computationally. More details on the activities of the two centers can be found in the laboratory summaries for Los Alamos and Oak Ridge in Section 3.

The Department of Energy has designated as Grand Challenges the following projects, which will be the initial targets for HPCC computing and communications technology.

- Global climate modeling involves the simulation of the coupled ocean-atmosphere global ecosystem.
- Groundwater transport involves the simulation of the movement of contaminants through soils and other porous media.

- Computer design of materials involves the *ab initio* simulation of the structure of novel materials and alloys.
- Computational chemistry involves the simulation of chemical processes that arise in environmental remediation and waste removal.
- Structural biology involves the simulation of protein folding and the analysis and interpretation of genomic sequencing and mapping.
- Petroleum reservoir modeling involves the simulation of flows through porous geologic media.
- Numerical tokamak involves the simulation of the physics of plasmas and the responses of materials inside a tokamak fusion reactor.
- Quantum chromodynamics involves the simulation of matter at the subnuclear level.
- Computational combustion dynamics involves the development and implementation of an adaptive mesh refinement methodology for simulating combustion phenomena in realistic geometries.

Together, the Department of Energy and the National Aeronautics and Space Administration have taken a major step toward the implementation of the National Research and Education Network. In August 1992 DOE and NASA selected Sprint to provide broad-band communications services that will upgrade the capabilities of three national research and data networks: the ESNet of DOE and the AEROnet and the Science Internet of NASA. This project will employ ATM (Asynchronous Transfer Mode) technology to enable these networks to move from the current transmission rate of 1.5 million bits per second to 622 million bits per second over the next few years.

Ongoing Work

There is ongoing work at the DOE laboratories that both supports the Grand Challenge projects outlined above and advances the state of the art in all four areas of the HPCC Program.

- HPCS. In the area of High-Performance Computing Systems, Fermi National Accelerator Laboratory has pioneered the development of hardware and software for clustering groups of inexpensive processors into centrally managed distributed parallel computers. In partnership with IBM and Merck, researchers at Fermilab have implemented parallelized crystallographic codes on clusters of workstations to simulate the behavior of proposed drug molecules with biological receptors and thereby reduce the excessive cost associated with traditional trial-and-error methods of drug design.
- ASTA. In the area of Advanced Software Technology and Algorithms, researchers at Lawrence Berkeley Laboratory and Lawrence Livermore and Los Alamos National Laboratories have developed data storage and data management software tools to handle the vast amounts of data generated by large-scale scientific and engineering projects, such as operating the Superconducting Supercollider, sequencing the human genome, and simulating complex phenomena like the global climate and the movement of environmental contaminants.
- NREN. To make effective use of the National Research and Education Network, high-speed network interfaces must be available for both supercomputers and workstations. Lawrence Berkeley Laboratory has developed a new type of workstation network interface that is inexpensive and yet delivers maximum performance.
- BRHR. In the area of basic research and human resources, applied mathematicians and computational scientists at Ames Laboratory have designed improved algorithms for reconstructing enhanced tomography scans from incomplete or noisy data and implemented them on massively parallel computers, thereby increasing the applicability of the scans to noninvasive medical diagnostics and to nondestructive evaluation of industrial materials. Ames Laboratory also collaborated with Oak Ridge and Sandia National Laboratories on an educational program called Adventures in Supercomputing. The program involves teachers and under-represented groups of high school students in computer activities and workshops and develops in the students an appreciation of science and computing and encourages them to pursue careers in science and technology.

Section 2 contains highlights of the first year of DOE's participation in the HPCC Program. The third and final section contains summaries of the projects supported by the Office of Scientific Computing.

DESCRIPTIONS OF FUNDING CODES FOR HPCC CATEGORIES

The B&R codes supply information for the four HPCC categories. Each code represents one subtopic. In FY92 funding codes for some categories were changed; in FY93 the conversion to the new codes will be completed. Thus, differences in codes will appear in this year's report only.

KC07 APPLIED MATHEMATICAL SCIENCES SUBPROGRAM

This subprogram supports research in and application of advanced mathematical, computational, and computer sciences to improve the ability of DOE to solve scientific problems that are critical to its mission and the nation.

KC0701 Mathematical, Computational, and Computer Sciences Research

This category includes the three B&R subcategories described below.

KC070101: Basic Research and Human Resources (BRHR)

This subcategory supports basic research and human resources development activities in mathematics, computational sciences, and computer sciences including improved mathematical representations of physical systems, algorithms, techniques for predicting the behavior of physical systems on advanced computers, and educational activities.

KC070102: High-Performance Computing Systems (HPCS) Research

This subcategory supports research and development on the underlying software, computer languages, and hardware architecture of HPCS, including the technological evaluation and characterization of new computer architectures and systems.

KC070103: Advanced Software Technology and Algorithms Research (ASTA)

This subcategory supports research and development of advanced software technology and algorithms that enable effective application of advanced computers to scientific problems that are critical to the DOE, including the early application of these technologies and algorithms, Grand Challenge collaborations, software tools environments, and computational techniques.

KC0702: Advanced Computation, Communications Research, and IRM (ACCRI)

This category supports research, development, and operations needed to apply advanced computational, communications, and mathematical techniques to scientific problems that are important for DOE and the nation by providing researchers supported by the Office of Energy Research with access to the advanced computational and mathematical capabilities and resources. It also supports the associated high-performance computing research centers, access centers, and communications infrastructures as well as the associated IRM functions required to enable this activity to succeed.

DESCRIPTION OF HPCC PROGRAM

The DOE HPCC program for FY93 will continue to build on the solid foundation of joint interagency, interdisciplinary, and private-sector collaborations that were established during FY92. The DOE HPCC Program seeks to enable the effective application of HPCC technology to scientific problems that are critical to the national energy strategy, other DOE mission programs, and the national interest. The program is, therefore, driven by three objectives: to advance the knowledge of mathematical, computational, and computer sciences needed to model and to provide better understanding of complex physical, chemical, and biological phenomena involved in energy production and storage systems; to advance the knowledge of high-performance communications techniques needed to support geographically distributed science collaborations and HPCC systems access; and to manage the DOE HPCRCs, including supercomputer facilities and the ESNet, as part of the overall DOE HPCC Program.

Because of this focus on applying HPCC technologies, technology transfer, and sharing between researchers in different disciplines and researchers and their counterparts in the US, commercial enterprises are a critical part of the DOE program. To ensure the effectiveness of these processes, the DOE program

- funds activities in cooperation with other programs that will be the eventual users of this technology;
- involves end users of the technology to the maximum extent possible in the initial and ongoing evaluation of projects; and
- developed, in conjunction with the Computer Systems Policy Project (CSPP), model language to address the legal aspects of Cooperative Research and Development Agreements (CRADAs) and streamline the process to facilitate the approval of CRADAs with industry partners.

HPCS

In HPCS research, the DOE initiated a performance measurement research program that has already played an integral part in the US–Japan trade negotiations. A DOE laboratory researcher developed, and was awarded a patent for, an innovative parallel systems benchmarking capability, SLALOM. And, the DOE continued two large university projects in parallel systems research (University of Illinois and New York University) and several ongoing technology evaluation projects.

ASTA

Support for Grand Challenges

The DOE conducted an interagency panel to evaluate proposals for Grand Challenge computational research projects in areas of importance to the DOE mission. These projects were selected via a peer review panel that included non-HPCC program managers and other HPCC agency participants and was chaired by a NASA program manager. The Grand Challenges selected are co-funded by other DOE programs and by industrial partners and include DOE lab, university, industry, and other HPCC agency participation. The six research areas selected are as follows:

- Computational Chemistry—to effectively parallelize important chemistry codes to study critical problems in the chemistry of halohydrocarbons, chemistry on and in clay minerals, and rational design of biodegrative enzymes for environmental remediation and future energy sources. This project involves co-funding from four industrial partners and the DOE Chemical Sciences Program.
- Computational Structural Biology—for research in computational methods of protein folding, parallel programming environments for biology, profile analysis for structural comparison and prediction, predictive models of cellular organelles, and large-scale sequence analysis. This project involves co-funding from the DOE Health and Environmental Research Program, an NSF Science and Technology Center, a biomedical industrial firm, and a university foundation.
- Mathematical Combustion Modeling—to develop algorithms for high resolution discretization techniques with adaptive methods, to enable effective parallelization of resulting combustion models, and to make these models useful to private sector and government scientists and engineers. This project is co-funded with the Applied Math Program.
- Lattice Gauge Calculations on Massively Parallel Machines—to develop algorithms to utilize full-scale HPCC systems to improve the quenched quark approximations and to further refine the standard model. This project is co-funded by the DOE and NSF high-energy physics programs.
- Petroleum Reservoir Modeling—for efficient use of HPCC systems to simulate flow through permeable media, especially oil flow simulations. The work is based on reservoir models developed at the University of Texas that are widely used by the oil industry and is co-funded by DOE Environmental Sciences, an industry consortium of twenty-six companies, two computer vendors (Cray Research, Inc. and DEC), NSF, and the state of Texas.
- Numerical Tokamak Project—to develop integrated particle and fluid plasma models on massively parallel machines as part of a multidisciplinary study of tokamak core fluctuations and transport of energy and particles. The measure of success will be the degree to which the observed scaling of fluctuation driven transport can be predicted. The project is co-funded by the DOE Fusion Energy Program and is an important component of the national energy strategy.

Software Tools and Components

The DOE initiated projects in the areas of computational catalysis studies, quantum chemistry and drug design, and molecular and materials sciences projects in addition to the Grand Challenges above and in cooperation with private sector firms.

High-Performance Computing Research Centers

The DOE created two HPCRCs—at Los Alamos National Laboratory and at Oak Ridge National Laboratory—involving industrial and university partnerships as well as co-funding from other DOE programs. These HPCRCs will conduct research in all four HPCC program component areas; will operate full-scale models of massively HPC systems prototypes installed at these HPCRCS during FY92 (i.e., a Thinking Machines Corp. CM-5 at LANL and an Intel Corp. Paragon at ORNL); will perform computational research in global climate research, environmental groundwater transport modeling, and materials sciences calculations; and will also serve as an intellectual center to evaluate the suitability of advanced computer architectures for Grand Challenge class problems, as well as to provide HPCS resources for computational research for all of the selected energy related to Grand Challenge research projects. The HPCRC at LANL is also a partner in a Science and Technology Consortium sponsored by NSF. The DOE also installed the Serial 1 Kendall Square computer system at the ORNL HPCRC to evaluate this new technology on energy and other applications as determined by peer review. This system is accessible on the ESNet, and therefore on the Internet and NREN.

In addition, the DOE installed a C-90 supercomputer system at NERSC, a supercomputer access center, as part of its base program activities to provide class VII computer technology in its access program to bridge the gap until a massively parallel computer can be introduced in a full production environment in late FY94. The C-90 will have 16 processors and will be set up to encourage parallel programming techniques to promote this transition. In conjunction with this activity, NERSC and Cray Research have entered into a CRADA to develop parallel software environments for this supercomputer and Cray's follow-on massively parallel processing systems, which will work in concert with the C-90.

These three centers at LANL, ORNL, and LLNL also serve as collaborative consortia, as set forth in the HPC Act of 1991.

NREN

The DOE conducted a competitive procurement and began upgrades of its Energy Sciences Network to 45 megabits service during FY92. This upgrade is being done in conjunction with NASA's Science Internet to merge existing Internet technologies with future high-performance communications technologies such as Switched Multimegabit Data Services frame and cell relay technologies. The DOE also incorporated its Office of Science and Technology Information into its ESNet sites. The DOE OSTI is the library and repository of all DOE science and technology publications and bibliographic information. OSTI also connects its bibliographic databases with those of other agencies and its international partners to provide DOE researchers and the private sector with the most recent and comprehensive energy information for use in all endeavors in research and education.

In the area of gigabit research, the DOE concentration is in optimizing and enabling packetization of workstation video and multimedia capabilities and in advancing local area network capabilities to handle data requirements of massively parallel systems for mass storage access, wide area network interfaces, etc. The DOE has initiated several projects in conjunction with DARPA and has held joint, coordinated program reviews of these projects with DARPA during FY92.

BRHR

In the area of education, the DOE has continued several excellent programs in its base program, including the DOE High School Science Student Honors Program; the Southwest Indian Polytechnic Institute, in which 30 Native Americans will compete using the MESA graphics program; and teachers' workshops to train teachers in HPCC technology and use and to train teachers to conduct such workshops themselves. Some of these workshops are co-sponsored by Cray Research, which donated a supercomputer to support this effort. Highlights for FY92 include the following:

- The NERSC education coordinator received a Public Service Award for her contributions in this important area.
- DOE initiated an "Adventures in Supercomputing" Program at three DOE labs to capture and cultivate the interest of women, minority, and disadvantaged precollege students in science, mathematics, and computing in conjunction with over 20 high schools in FY92. This program includes summer teacher fellowships at DOE labs and is coordinated with many other high school programs, such as the New Mexico Supercomputer Challenge, the Alabama Precollege Supercomputing and SuperQuest Programs, and the Colorado State computational science program.
- The DOE has drafted and made available on the ESNet and interagency interim NREN a syllabus for teaching computational science at the graduate level. This work involved 24 authors from 10 disciplines. It is in electronic format and incorporates graphics that use Apple Macintosh and IBM RS 6000 technologies and allow one to execute examples on IBM, Intel, Cray, Thinking Machines, and Kendall Square systems. It is electronically available and maintained at Vanderbilt University.

In the area of applied mathematics, the DOE and NSF have co-funded the Geometry Science and Technology Center at the University of Minnesota to apply new HPCC technology to traditional mathematics problems and educational programs. DOE has co-funded a Grand Challenge project in mathematical combustion modeling, and has initiated new efforts in complex nonlinear behavior, which underlies many natural phenomena, and in graph and group theories related to discrete phenomena and topology for application to genome sequencing and protein structure and folding.

PROGRAM MANAGEMENT

The DOE has assigned the Office of Scientific Computing lead responsibility for DOE-wide participation in the HPCC Initiative. During FY92, the DOE HPCC Program has made many significant accomplishments in HPCC projects and changes to the existing DOE Applied Math Program in concert with the President's Initiative in HPCC and with its 5-year goals. In addition, the DOE is committed to using, wherever feasible, technologies and research conducted by other agencies to support its application of HPCC technologies.

FEDERAL REGISTER ANNOUNCEMENT

A Special Research Grant Program Notice in the Federal Register announces the period during which the Office of Energy Research will accept formal applications for grants for the next fiscal year. The notice, reproduced below, states the amount of appropriated funds, describes the major components of the HPCC program, and provides application and award information. The Office of Energy Research provides to requestors copies of the application kit (with instructions), 10 CFR Part 605, and reports on research funded during prior fiscal years.

US DEPARTMENT OF ENERGY
OFFICE OF ENERGY RESEARCH

SPECIAL RESEARCH GRANT PROGRAM NOTICE 92-20: HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS

Agency: US Department of Energy

Action: Notice inviting grant applications

Summary: The Scientific Computing Staff of the Office of Energy Research (ER), US Department of Energy (DOE), hereby announces its interest in receiving applications for Special Research Grants supporting DOE's role in the President's High-Performance Computing and Communications (HPCC) Initiative. The 5 year federal HPCC Program was announced by Dr. D. Allan Bromley, Director, Office of Science and Technology Policy, in February 1991. In December, the 102nd Congress passed the "High-Performance Computing Act of 1991," which provides for a coordinated federal program to ensure continued United States leadership in high-performance computing. The DOE has an integral and broad program within the eight-agency HPCC Initiative. The primary goals of the DOE HPCC Program are: (1) to extend US technological leadership in high-performance computing and computer communications; (2) to improve US productivity and industrial competitiveness by making high-performance computing and network technologies an integral part of the design and production process; and (3) to provide wide dissemination and application of the advances in these technologies to both speed the pace of innovation and serve the national economy, security, and education.

The DOE program will approach these goals by (1) supporting research and development to solve important scientific and technical challenges; (2) reducing the uncertainties in industrial research and development through increased cooperation between government, industry, and universities and by continued use of government and government-funded facilities as a prototype user of early commercial HPCC products; (3) supporting the underlying research, network, and computational infrastructures on which US high-performance computing technology is based; and (4) supporting the US human resource base to meet the needs of industry, universities, and government.

This notice requests applications for grants to support research in these major components of the HPCC Program:

1. High-Performance Computing Systems (HPCS)—research to advance the capabilities of future generations of computing systems and to evaluate advanced prototype systems;
2. Advanced Software Technology and Algorithms (ASTA)—software support for the computational Grand Challenges by research and development of software tools, components, and computational techniques and by the establishment of High-Performance Computing Research Centers (HPCRCs);
3. National Research and Education Network (NREN)—research and development on very-high-speed digital communications (gigabits) and participation in the interagency interim NREN; and
4. Basic Research and Human Resources (BRHR)—research participation and training, educational infrastructure, education, and curriculum development.

Collaborative research among investigators at universities, industrial firms, and DOE national laboratories is encouraged. Advanced software technology and algorithms in support of the DOE energy-related computational Grand Challenges will be emphasized.

Dates: The DOE HPCC initiative is a 5-year program. To permit timely consideration of awards in Fiscal Year 1993, formal applications submitted in response to this notice must be received by November 4, 1992. Earlier submission is encouraged.

Addresses: Formal applications referencing Program Notice 92-20 sent by US mail should be forwarded to: US Department of Energy, Office of Energy Research, Division of Acquisition and Assistance Management, ER-64, Washington, D.C. 20585, ATTN: Program Notice 92-20. The following address must be used when submitting applications by US Postal Service express or any commercial mail delivery service or when hand-carried by the applicant: US Department of Energy, Office of Energy Research, Division of Acquisition and Assistance Management, ER-64/GTN, 19901 Germantown Road, Germantown, MD 20874.

For Further Information Contact: The Scientific Computing Staff, Office of Energy Research, ER-7/GTN, US Department of Energy, Washington, D.C. 20585, (301) 903-5800.

Supplementary Information: The DOE HPCC Program is described in a report, "The DOE Program Component of the Federal High-Performance Computing and Communications Program," dated June 1991 (DOE/ER-0489P). The FY93 federal program is summarized in "Grand Challenges 1993: High-Performance Computing and Communications—A Supplement to the President's FY 1993 Budget." These reports can be obtained by calling (301) 903-5800.

DOE participates in all four major components of the federal HPCC Program as described below.

1. High-Performance Computing Systems (HPCS). In the area of HPCS, DOE will be an early customer of small versions of systems with advanced architectures and will evaluate these systems on energy-related applications. It should be noted that the primary technology development for HPCS of the HPCC Program will be managed and funded by the Defense Advanced Research Projects Agency (DARPA). DOE will consider cooperative development projects of advanced systems involving its national laboratories, universities, and vendors. DOE will support research and development associated with evaluating the effectiveness of new parallel computing systems.

2. Advanced Software Technology and Algorithms (ASTA). The ASTA component is strongly emphasized in DOE's HPCC Program. The effort will provide support for Grand Challenge and other computational collaborations pertinent to the DOE missions; software components and tools for high-performance systems and prototype computational science programming environments that are portable and conform to standards or *de facto* standards; and high-performance computing research centers to facilitate research on parallel machines to enhance the effectiveness of applications and programming environments. In addition to these three elements, described in detail below, the HPCC Program will support research and evaluation aimed at improving the integration of advanced DOE research in computational science into commercial products and processes with significant economic impact for the country. Support is also available for pertinent research on the needs of the computational community.

2.1. The Grand Challenge and computational projects subcomponent is directed at enabling major advances in science and engineering of critical importance to DOE by providing computational resources—time on high-performance computing systems and advanced, effective software. These projects will be focused on science or engineering applications and are expected to involve multidisciplinary teams of scientists and engineers, along with computational researchers and computer scientists from academia, governmental institutions, and industry.

Goals of the HPCC Initiative include accelerating the pace of innovation and enhancing US competitiveness in the marketplace; these provide the emphasis in this subcomponent. To insure the purposes of the initiative are fulfilled, applications for grants in this subcomponent should address the following questions:

- Fundamental Significance. What fundamental science or engineering problem is addressed? What is the expected economic, societal, and scientific impact?
- International Competitiveness. What contribution to the nation's productivity/economy and international competitiveness will the proposed project have? Industrial involvement is the most evident contribution, and partnerships with industrial users of HPCC are strongly encouraged.
- DOE Mission. What is the contribution of the proposed project to the missions of DOE? If a project pertains to a particular DOE science office or a technology office, then the goals of the project should be appropriate to the missions of that office as well as to the missions of the office of the Scientific Computing Staff.
- Enabling Technologies. Will the expected progress in software and hardware technologies enable a substantial advance in the project in the next year or two?
- Interdisciplinary Approach. Is an interdisciplinary approach involving scientists, engineers, mathematicians and computer or computational scientists appropriate? Only under rare circumstances would an HPCC computational project not incorporate an interdisciplinary approach.
- Support Leverage. What is the support for the project from present and pending sponsoring partners—DOE program offices, other federal agencies, private or non-profit institutions, and industrial firms?
- Technology Leverage. What probable advances in enabling software or hardware technologies developed in this project will benefit other computational projects?

2.2. The software technology subcomponent of the DOE program is focused on improving the effectiveness of computational scientists and engineers who are working on Grand Challenge class problems of critical importance to DOE and the nation. This technology falls into three general categories: (1) technology to assist in constructing, maintaining, and monitoring scientific software in high-performance computing environments; (2) technology for understanding and managing the output of scientific software; and (3) software infrastructure, including compilers, libraries, and computational techniques.

Also, applications should include the following additional information:

- Their potential, relative to other efforts supported by DOE and other HPCC agencies for producing technology that is needed and desired by computational scientists and engineers.

- Their potential for producing technology that can interoperate with other technologies through the use of standards and open, clear, rich interface definitions.
- Their potential for producing technology that has a clear path leading to commercial products either through demonstrated prior success or partnership with commercial concerns.
- The quality and level of involvement of computational scientists from government, industry, and academia who are working on Grand Challenge class problems or who are involved as users and evaluators or advisors for the software technology being proposed.

2.3. Because of the intricate nature of HPCRC applications and uncertainties in the budget and scope of the DOE HPCC Program in FY93, please contact the Scientific Computing Staff before compiling such an application. Following is a description of the characteristics DOE expects of an HPCRC.

The DOE's HPCRCs are to serve as intellectual homes for leading-edge research that will enable the solution of Grand Challenge problems. The activities at the HPCRCs are designed to support interdisciplinary and inter-institutional collaborations. The HPCRCs will house interdisciplinary teams of researchers drawn from science and engineering applications disciplines, computational mathematics, and computer science. These teams focus on application development to harness the power of advanced architecture computers for the solution of computationally intensive, large-scale scientific and engineering problems. Multidisciplinary research also will be emphasized. Computational science and engineering should transcend traditional disciplinary boundaries and present opportunities for novel interactions and synergies. The HPCRCs are expected to exploit these opportunities.

The HPCRCs provide a natural setting for a variety of supporting research activities, such as computer performance evaluation, visualization, software environments for computational science and engineering, gigabit networking research, high-speed data storage and archiving, and the development of new solvers and better software libraries. The HPCRCs are expected to play an active role in maintaining a dialogue with industry, universities, and other laboratories and centers in order to maximize the dissemination of information, promote and support technology commercialization, and avoid unnecessary duplication of effort.

The HPCRCs are designed to play an important role in computational science education. This aspect of their mission is expected to go well beyond the usual training courses in computing and to embrace such issues as computational science curriculum design at the university level and introductory and motivational programs and material for precollege students and teachers. These activities require strong collaboration with the academic community.

While they are not production computing centers, the HPCRCs house and vigorously exploit full-scale versions of prototype advanced architecture computers. In order for remotely located collaborators to fully participate in the research activities of the HPCRCs, each center will become a major node on a NREN backbone.

Because high-performance computing and its applications are expected to grow rapidly and experience rapid changes in technology, DOE's commitment to any particular HPCRC will be of limited duration. However, new HPCRCs may be created to exploit new computing technologies and future Grand Challenge applications.

3. The National Research and Education Network (NREN). DOE will participate in the cooperative interagency NREN. The Energy Sciences Network (ESNet) will be incorporated into NREN to provide quality network access to the energy research facilities by research and education communities. It should be noted, however, that broad community access to NREN will be supported by the National Science Foundation (NSF) through the NREN component of the federal HPCC Program. ESNet will maintain compatibility and will be upgraded in concert with NREN. Applications are invited to develop gigabit network support technology for DOE projects distributed across multiple energy research centers at the national laboratories and universities. Primary coordination and funding for gigabit research in the federal HPCC Program will be done by DARPA.

4. Basic Research and Human Resources (BRHR). DOE's BRHR activities will include: stimulating research and education in computational science; expanding training programs for high school teachers and college students in computing techniques; initiation of high school supercomputer programs; and provision of fellowships in computational science with internship at national laboratories.

Application and Award Information

Information about submission of applications, eligibility, limitations, evaluation and selection processes, and other policies and procedures may be found in the Application and Guide for the Special Research Grant Program and in 10 CFR Part 605. The application kit and guide and copies of 10 CFR Part 605 are available from the Office of Energy Research, Scientific Computing Staff, ER-7, Washington, D.C. 20585. Instructions for preparation of an application are included in the application kit. However, the project description should not exceed 25 double-spaced pages. Lengthy appendices are discouraged. Telephone requests may be made by calling (301) 903-5800. The Catalog of Federal Domestic Assistance number for this program is 81.049.

Subject to availability of appropriated FY93 funds, approximately $3,000,000 will be available for award. The allocation of funds will depend upon the number and quality of applications received. Grant awards will generally be for a three-year period, be funded one year at a time, and may range from $50,000 to $850,000.

Issued in Washington, D.C., on August 13, 1992.

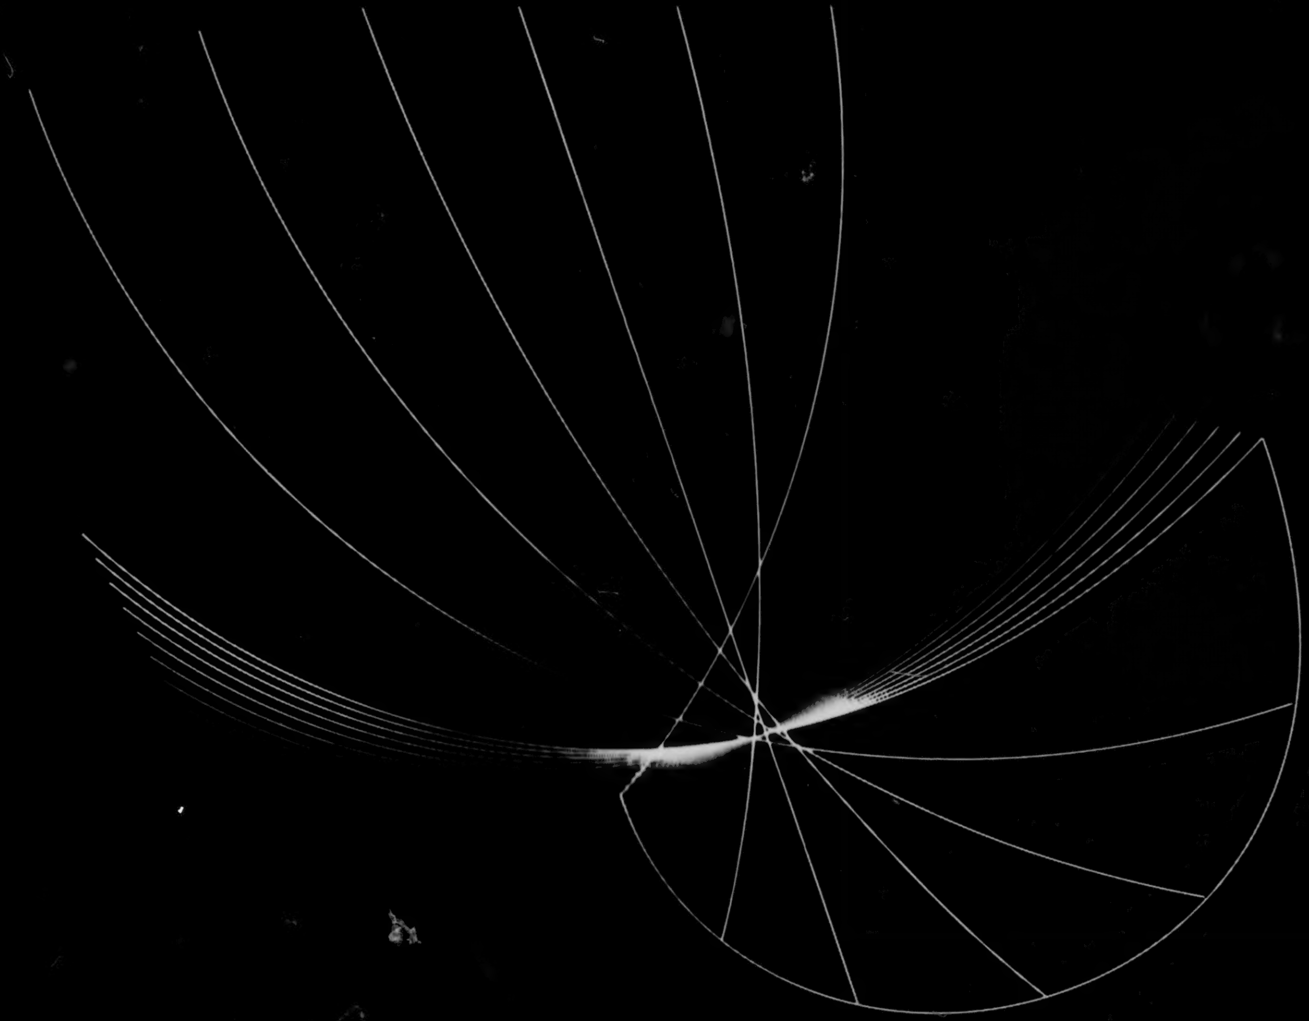

Section 2
Highlights

ADVENTURES IN SUPERCOMPUTING

The Ames Laboratory, Sandia National Laboratories at Albuquerque, and Oak Ridge National Laboratory have developed an educational program designed to interest high school women and minorities in computational science by providing access to supercomputers at the three laboratories. Earlier this year, five Iowa high schools were selected by a committee made up of Iowa Department of Education officials and Ames Laboratory personnel to take part in this program. The high schools will be provided with four computers (Macintosh si model), software, and a 56-KB network connection to Ames Laboratory and the nCube in the Scalable Computing Laboratory at the Ames Laboratory. During a two-week summer institute, the high school teachers in New Mexico, Tennessee, and Iowa received training on a supercomputing curriculum developed at the University of Alabama at Huntsville.

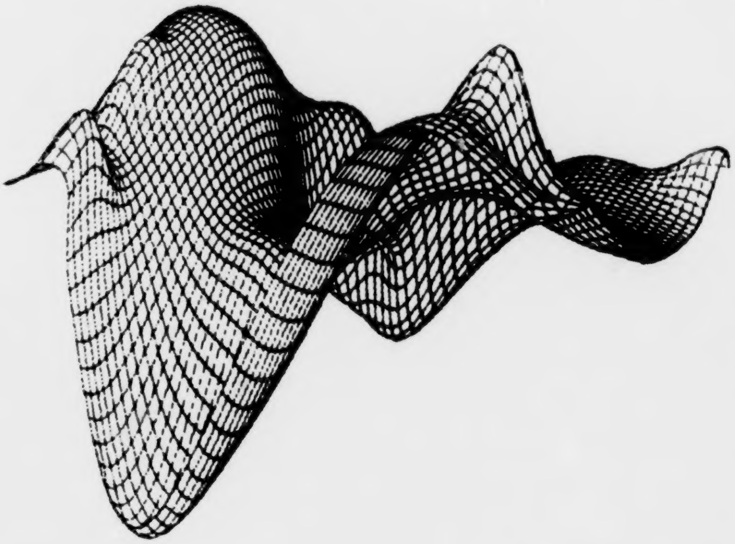

Argonne National Laboratory

ADVANCES IN CARDIOVASCULAR MODELING ENHANCED BY NEW AUTOMATIC DIFFERENTIATION TOOLS

A software tool developed at Argonne National Laboratory has provided new insights in studies under way by researchers at the University of Pittsburgh Department of Surgery in collaboration with numerical analysts at the National Institute of Standards and Technology (NIST). These studies involve mathematically modeling the biomechanical implications of disease on cardiovascular tissue. The software tool, called ADIFOR (Automatic Differentiation in FORTRAN), calculates the derivatives required for this modeling effort more accurately and efficiently than conventional methods.

The University of Pittsburgh researchers hypothesize that the elastic properties, and thus the mechanical function, of cardiovascular tissue are altered by disease, and they have proposed a method to correlate cardiovascular disease with the change in these elastic properties. The goal is to eventually enable physicians to determine disease level (or severity) noninvasively.

The researchers have developed a mathematical model to predict the mechanical function of cardiovascular tissue and are now working with NIST analysts to find the parameters of this model that produce the "best fit" to their experimental data. The fitting procedure requires values for the derivatives of the model with respect to each of the parameters, however, and because of the model's complexity, traditional divided-difference approximations to these derivatives are not adequate. ADIFOR offers a solution.

By providing reliable derivatives, ADIFOR has enabled the researchers to move from the problem of determining reliable divided-difference approximations to the more interesting issues of parameter estimation. The results obtained with the ADIFOR-generated derivatives have identified problems with the current model, and the researchers are now working on improvements. Because ADIFOR derivatives can be generated with little effort, the scientists will be able to easily integrate the derivatives for new or modified models into their existing data-fitting framework.

Experimental results show that ADIFOR's approach to automatic differentiation can dramatically reduce the time required to compute derivatives, thus reducing the overall computation time for fitting models to data. NIST analysts have also determined a secondary benefit of using ADIFOR-generated derivatives with least-squares data-fitting procedures. They have found that the improved accuracy of these derivatives as compared to the traditional forward divided-difference derivatives can produce a significant saving in the time required for the linear algebra calculations within the fitting procedures. In one camera calibration problem analyzed at NIST, for example, the solution found using ADIFOR-generated derivatives involved significantly fewer internal adjustment steps than were required when divided-difference derivatives were used. The result was an overall 60 percent reduction in execution time. In another NIST analysis of microwave resonant cavity data, the reduction was 80 percent. For large models, it may be possible to reduce execution time by orders of magnitude.

According to NIST researchers, ADIFORS benefits are clear: "Automatic differentiation derivatives improve performance and are easy to generate, reliable, cost-effective, and sometimes absolutely necessary for success."

Fermi National Laboratory

PARALLEL WORKSTATION FARMS: EXPRESS PATH TO PARALLEL COMPUTING IN INDUSTRIAL RESEARCH

Project Summary

Fermi National Accelerator Laboratory has pioneered the development of hardware and software systems for clustering groups of inexpensive processors into centrally managed distributed parallel computers called "farms." These systems were originally created to meet Fermilab's needs in High Energy Physics (HEP). A collaboration with IBM has been underway to refine the farm approach to solving large computational problems with cost-effective and extensible technology. The addition of Merck & Co., Inc. to this team and the installation of a farm at their location guarantees the application of farm computing technology in another critical and industrially based scientific domain, Computer-Assisted Drug Discovery. Characterized by moving innovations quickly into the production stream and onto the benchtop, Merck is challenged to use the newest hardware and software to model the properties and interaction of proposed drug molecules with biological receptors. They have parallelized crystallographic codes; built client/server heterogeneous environments for molecular modeling; and performed experiments with distributed parallel computing on workstations.

Fermilab currently operates a 140-node production farm of IBM RISC System/6000s, interconnected with Ethernet and configured with file and tape servers. Fermilab's Cooperative Processes Software is a package of software tools that enables a user to partition jobs into logical classes of processes and then distribute and manage them across multiple CPUs. A batch system that manages the job execution, resource allocation, and system administrative tasks has been added. This is a feature necessary for large multiuser, multiplatform production environments, such as Merck's. We are planning enhancements to the existing farm environment that facilitate the use of large databases, such as chemical information-management systems, and provide intelligent scheduling of resources and more efficient distributed data handling. These are features that will be required to deliver the "look and feel" of large data centers to increasingly demanding users.

Report

The configuration of the 16-node farm (DOE and IBM co-funded) for development work at Fermilab and Merck is underway at Fermilab. IBM has been exploring the use of Fiber Channel to interconnect subsets of processors in the farm at high speed, in anticipation of making that enhancement to the Fermilab-Merck farms. IBM has lent Fermilab prototype new technology; CPS has been successfully ported to it and jobs executed. Merck has procured 4 high-end IBM RISC System 6000 workstations as the first phase of their permanent farm. These workstations have been connected via a Network Systems Corporation router using fiber optics into a high-speed local network, together with their CRAY YMP.

Merck has installed and tested an earlier version of CPS and will receive the enhanced version just completed by Fermilab during a training residence at Fermilab in January 1993. They have also selected and begun to parallelize two key computational chemistry codes for execution on the farm. Fermilab is completing the CPS documentation for public release to coincide with the visit from Merck. Fermilab has added visual tools for console and system management to CPS that will facilitate the optimization of system performance.

Fermilab's "farm" of clustered groups of inexpensive processors.

IBM and Fermilab are developing concepts for a new class of applications involving group analysis of large quantities of shared data. This will apply tightly coupled MPP systems to act as "analysis servers." Merck will be invited to share in this effort because of its potential relevance to much of their drug design computation.

Merck has already begun a parallel development of an intelligent scheduler for a heterogeneous UNIX environment with IBM. They will investigate the utility of Fermilab's batch system in their environment and will then decide on the direction of their development efforts.

EQUILIBRIUM-FREE SURFACE INTERFACES

Our mathematical and computational studies have uncovered a class of "exotic" containers that can be constructed for any fluid/container materials pair. These containers, which are rotationally symmetric, have striking "symmetry-breaking" properties. One such container, which differs only locally near the mid-height from a circular cylinder by having a toroidal-like bulge, when half filled with fluid admits an entire continuum of distinct symmetric equilibrium free surfaces, including the horizontal planar surface. All of these surfaces make the same angle of contact with the container, enclose the same volume of liquid, and have the same mechanical energy. This behavior is in striking contrast with the familiar one of a cylindrical container, for which under the same conditions only a single surface is possible. Although the symmetric surfaces in the exotic container are equilibrium ones, they are unstable in that do not provide a local minimum for the energy. It can be shown furthermore that there is a stable, minimizing surface but it has to be asymmetric, even though the container itself is symmetric.

Exotic containers can be characterized for any contact angle and strength of gravity field. They become exceedingly small in scale, however, and effects correspondingly are very difficult to observe, except under microgravity conditions. Present-day possibilities for experimentation in the extended low-gravity environment of orbiting space vehicles allow the mathematical results to be tested experimentally. We are taking full advantage of these unique opportunities. In collaboration with NASA engineers, we have carried out drop-tower experiments at the NASA Lewis Zero Gravity Facility. Based on the results and our computer simulations, an experiment was designed to be carried out in space and was flown on the first NASA United States Microgravity Laboratory (USML-1) Space Shuttle flight in June and July 1992. The goals of the experimental aspects of our study are both to substantiate the mathematical results and to test the limitations of the underlying classical Young-Laplace theory of capillarity.

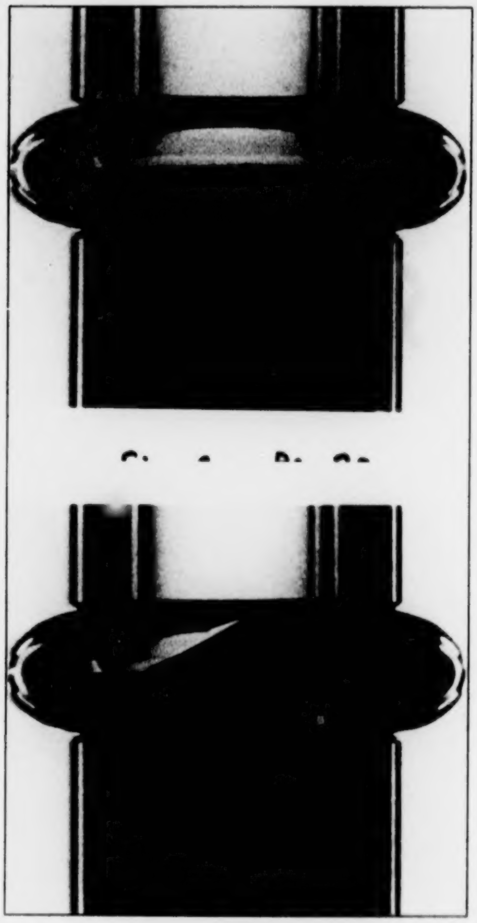

During a drop-tower experiment, fluid reorients in an "exotic" container from an initial configuration having a flat horizontal interface to a non-symmetric one during weightlessness (zero gravity), in agreement with mathematical theory.

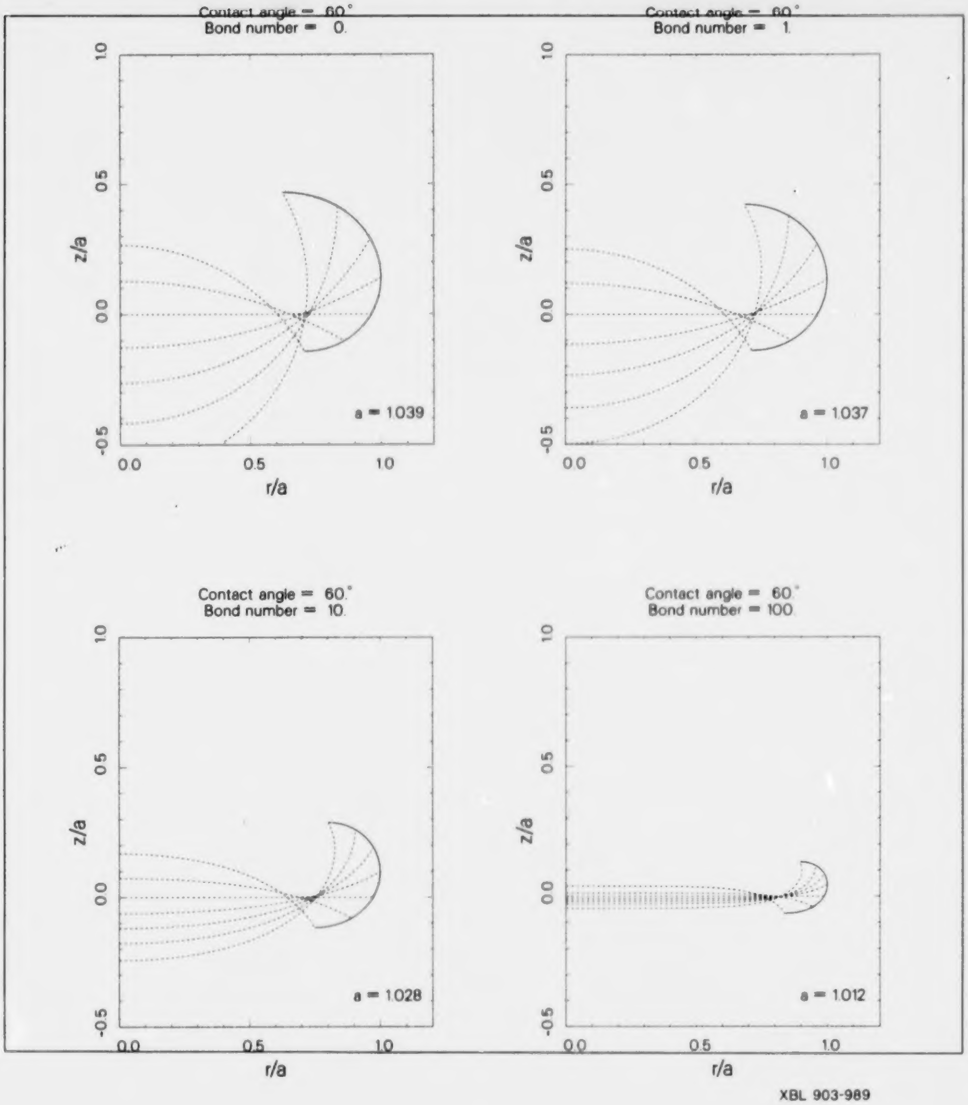

Exotic container meridian (solid curve) for contact angle 60 degrees, for Bond Number (dimensionless parameter proportional to the strength of the gravity field) 0, 1, 10, and 100. The dashed curves are meridians of members of the continuum of symmetric free surface interfaces all making the same contact angle with the container, enclosing the same volume of liquid, and having the same mechanical energy.

Lawrence Livermore National Laboratory

ADAPTIVE METHODS FOR COMPUTATIONAL FLUID DYNAMICS

Fluid dynamics problems are often characterized by variations on many length scales and by complex spatial and temporal behavior. In combustion, the energy release takes place in a region, such as a flame front, whose thickness is much smaller than the typical length scales of the large-scale motion of the fuel and air. In flow past immersed bodies, the forces on the body are completely determined by the thin viscous boundary layers near the body and by the behavior of the spatially confined wake behind it. Fluid flows arising in realistic engineering devices are sensitive to the details of the shape of such devices. Over the last ten years, a research effort in Computational Fluid Dynamics (CFD) funded by the Applied Mathematical Sciences program has been aimed at developing numerical methods that accurately and efficiently represent these effects. The basis for these numerical methods are high-resolution finite difference schemes that build the dominant physical behavior of the partial differential equations into the algorithm. We are combining these discretization techniques with adaptive methods that focus computational effort to obtain optimal performance on target applications.

These adaptive techniques include local adaptive mesh refinement (AMR) and front tracking. In local adaptive mesh refinement grid points in the finite difference calculation are dynamically inserted and removed to maintain a desired level of accuracy and resolve important time-dependent features in the solution to a problem. Our approach organizes refined grid points into a hierarchy of subgrids that are locally topologically rectangular. This strategy leads to minimal data management overhead, small storage requirements, and effectively exploits our high-resolution finite difference methodology. The other type of adaptive technique we are developing is front tracking, in which specialized algorithms are used to explicitly model the propagation of thin fronts, such as shock waves, flame fronts or material interfaces through the computational domain. Our approach is based on a volume of fluid representation in which the front is represented in terms of volume fractions. Jump relations are used to reconstruct the front and are applied as boundary conditions on either side of the front to determine its time evolution. This type of approach localizes geometry, easily accommodates changes in the topology of the tracked front, and provides a viable strategy in both two and three dimensions. The combination of these techniques typically reduces computational cost by a factor of 10–100 for a given level of accuracy when compared to conventional fixed grid methods.

An example of results obtained using these techniques is displayed in the figure.

THE NATIONAL STORAGE LABORATORY

The National Storage Laboratory (NSL) is a collaborative, high-performance storage project organized to investigate and commercialize technologies that promise to remove storage system bottlenecks and provide new needed storage system functionality. The NSL will provide new storage system functionality to help meet the storage requirements of critical DOE, other government, and US industry applications; will feed back to commercial vendors new methods of effectively utilizing storage hardware and software; and will guide the efforts of national storage system standards groups.

Large scientific, engineering, and commercial applications are straining storage and networking facilities, a condition compounded by new supercomputers, massively parallel processors, and high-performance workstations. The NSL focuses on bringing large-scale hierarchical storage system performance and functionality into better balance with modern processing, networking capabilities, and application requirements, so that high-performance computing can realize its full potential to support large simulation and modeling applications. The major goal of this project is to leverage DOE laboratory experience, expertise, and applications with that of industry to accelerate the development of US technology for local and nationwide storage system architectures and functionality and to facilitate and speed the availability of improved systems to DOE, other government agencies, and US industry.

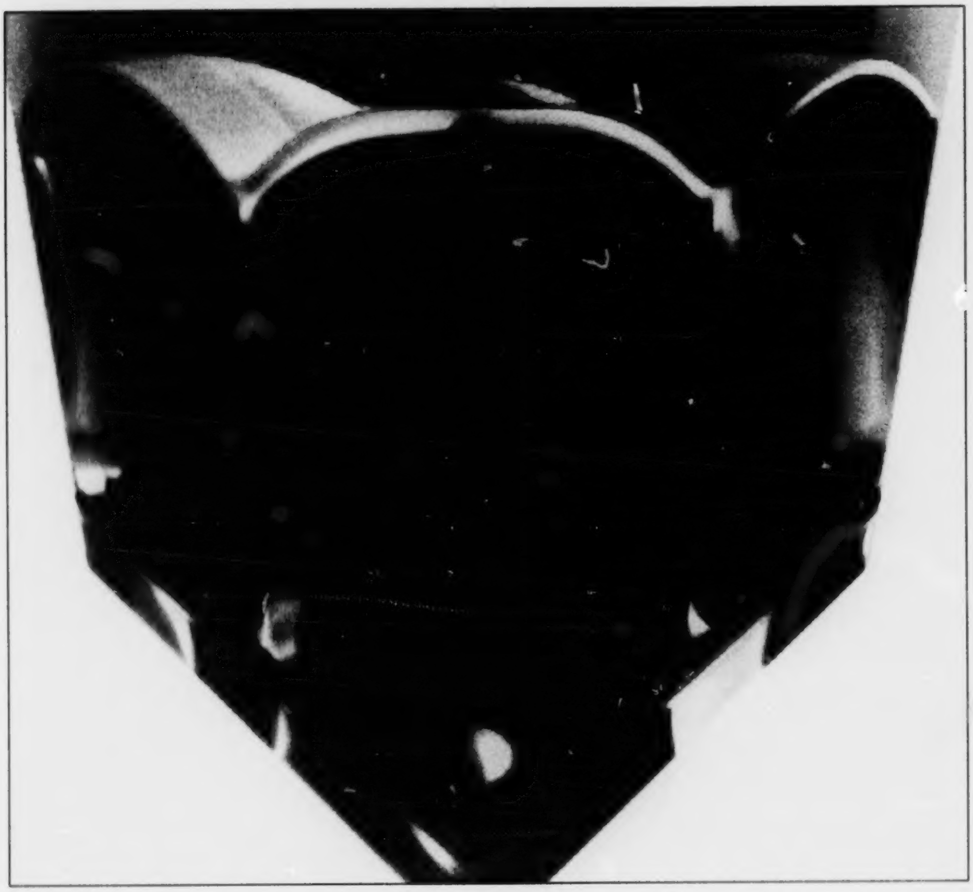

This figure illustrates the use of adaptive computational techniques to model complex three-dimensional flows in realistic geometries. We show the time evolution of jets formed by high pressure gas flowing through the holes in a baffle plate into a cylindrical chamber (only one-quarter of the cylinder is displayed). The jets emerge into the chamber above the plate from the four circular holes in the plate, driving a compression wave (light) ahead of the jets, followed by the expanding jet material (dark). In this example, we use a specialized interface modeling technique to represent the cylinder and the plate as reflecting surfaces embedded in a Cartesian grid. Local adaptive mesh refinement techniques are used to resolve the complicated fluid dynamics of the jets.

The NSL is located within the National Energy Research Supercomputer Center (NERSC) at Lawrence Livermore National Laboratory (LLNL) and will be integrated into the DOE Energy Sciences network (ESnet) and LLNL's Open LabNet environment for testing and demonstration with real applications.

Three initial LLNL applications have been chosen to help set priorities, test, and demonstrate the expected increase in productivity possible with the NSL system architecture: (1) the Program for Climate Model Diagnosis and Intercomparison (PCMDI), (2) Magnetic Fusion Energy (MFE) modeling and experimentation, and (3) digital storage and replay of computer-generated scientific movies.

The approach of the NSL is based on the premise that no single company, government laboratory, or research organization has the ability to confront all the system-level issues that must be resolved before there is significant advancement in storage systems technology. In an attempt to find solutions to these kinds of storage problems, LLNL has joined with industry and other laboratories and agencies in this collaborative project.

The initial participants besides LLNL are the IBM Federal Systems Company, IBM ADSTAR, Ampex Corporation, General Atomics' DISCOS Division, Maximum Strategy Incorporated, Network Systems Corporation, Zitel Corporation, IGM, PsiTech, CHI Systems, Cray Research, and the four National Science Foundation Supercomputer Centers. The industrial partners are initially providing over two million dollars' worth of equipment and over three million dollars' worth of personnel over the life of the project. Los Alamos National Laboratory, the Sandia National Laboratories, and other DOE laboratories will be joining the collaboration. Additional growth in the collaboration is expected.

The NSL has sparked worldwide interest because of the importance of the problems being addressed and strength of the collaboration. The first commercial fruit of the project was IBM's announcement in November 1992 of the NSL-UniTree storage system that supports the architecture in the figure and new NSL-developed functionality. IBM and General Atomics DISCOS Division have signed cross licensing agreements that will make the system available on many vendor's systems.

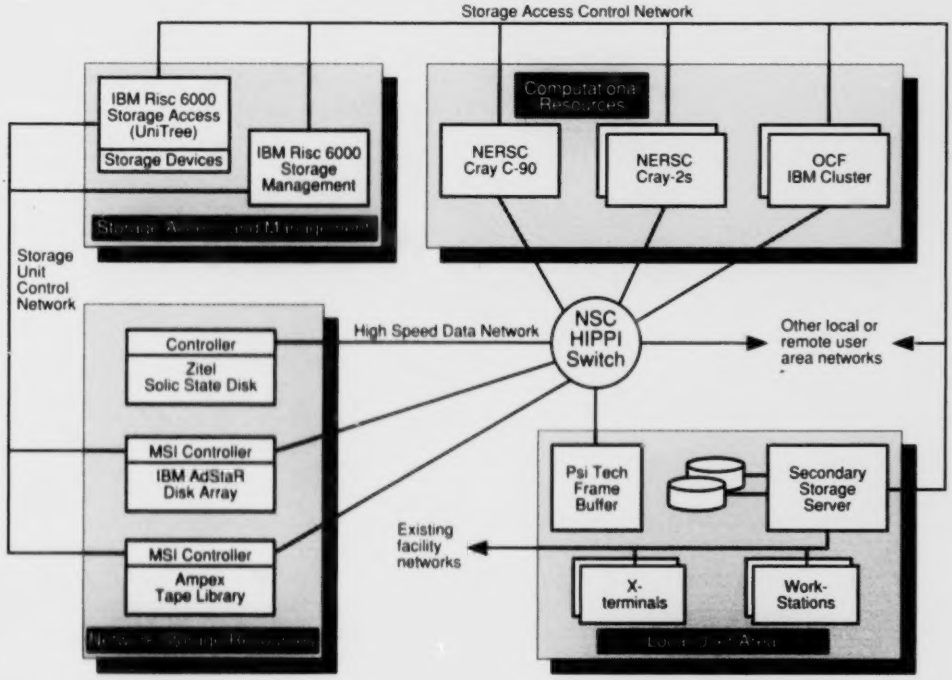

Storage Access Control Network
IBM Risc 6000 Storage Access (UniTree)
Storage Devices
IBM Risc 6000 Storage Management
Storage Access and Management
Computational Resources
NERSC Cray C-90
NERSC Cray-2s
OCF IBM Cluster
Storage Unit Control Network
High Speed Data Network
NSC HIPPI Switch
Other local or remote user area networks
Controller
Zitel Solic State Disk
MSI Controller
IBM AdStaR Disk Array
MSI Controller
Ampex Tape Library
Network Storage Resources
Psi Tech Frame Buffer
Secondary Storage Server
Existing facility networks
X-terminals
Work-Stations
Local User Area

The National Storage Laboratory configuration.

VISUALIZATION FOR GLOBAL CLIMATE MODELING

This effort focuses on new techniques for representing the vast amount of information generated and requiring analysis from the global climate models being studied by the Program for Climate Model Diagnosis and Intercomparison (PCMDI). Research into better representations of clouds, wind fields, and heating properties were identified as key areas of needed improvement for the early stages of this project.

LLNL has extended the state of the art in representing 3-D density distributions similar to those being used for medical CAT and MRI analysis to the curvilinear topology of the global climate models, as well as the topology generated by wrapping this curvilinear topology around a sphere. Work continues on a flexible framework that will permit extension of these techniques to handle several 3-D scalar fields or vector fields. This will enable researchers to explore the relationships between the clouds and radiation or global heating terms. Algorithms to texture the clouds and then advect those textures via the wind velocity fields were developed and presented. On a first pass, this provided realistic clouds advecting with the wind field. Work on controlling the textures and advection to better visualize the 3-D wind fields is progressing. Turbulent vector fields can now be represented using high-frequency textures. The technique has been applied to 3-D wind fields. This technique allows for the 3-D distribution of cloud densities to be integrated with the vector fields. Various controls are offered on both the representation of the scalar field and the vector field. Work is also progressing on this algorithm to incorporate the curvilinear terrain topology and provide the proper set of controls for the researchers. Initial reactions to all of these tools in the global climate modeling community are positive. Work remains to transfer these algorithms for day-to-day use by this community.

Los Alamos National Laboratory

The New Mexico Supercomputing Challenge is an academic-year-long competition that was initiated in 1990. Now in its third year, the Challenge is a program that allows high school students to do computational science projects in teams using high-performance computers. The Challenge is open to all New Mexico students on a nonselective basis. The purpose of the Challenge is to expose students and teachers to computational subjects and experiences that they might otherwise not have. Students receive training, continuous coaching, technical support, and an opportunity to win university scholarships and savings bonds to help further their education. Everyone who participates in the Challenge is a winner. The Challenge serves as a model that can be enhanced and adapted for other states. The Challenge is sponsored by a partnership of business, universities, and national laboratories.

The demand for the Challenge increased almost two-fold from the first to the second year:

	1990–91	1991–92
Students	235	419
Teams	65	112
Schools	40	51
Teachers	55	91

We expect 700–1000 student participants during this third year. Regarding ethnic distribution of student participants, 45% were minorities in 1991–92. More than 50% of the participating schools are in rural areas. In addition, in 1992 students used a much larger amount of computer time: 446 Cray CPU-hours in 91–92 compared with 80 hours in 90–91 and 28 connection machine hours in 91–92 compared with 19.5 hours in 90–91.

The winners of the 1992 Challenge demonstrate their work to one of the judges.

THE COMPUTATIONAL SCIENCE WORKSHOP

The Computational Science Workshop is specifically designed to enable researchers from throughout the nation to apply innovative technology to existing research projects. The workshop begins with a three-week intensive seminar, and in the following six weeks there are approximately 60 additional hours of seminars. Seminar topics are presented in three areas: (1) Methodologies (Numerical Methods, Parallel Algorithms), (2) Environments (Architectures, Visualization, Performance Languages), and (3) Applications (Grand Challenge Applications).

The Summer 1992 Workshop had 42 participants, 40% of whom were from academia, 40% from national and military laboratories, and 19% from industry. The benefits of the workshop were not only acquisition of new knowledge and abilities related to computational science, but also establishment of new collaborations. Some examples of the research results produced during the workshop included Computation and Visualization of 2-D Invariant Manifolds, Examination of Parallelism in Eigenvalue Solvers using Dataflow, Supercomputer Application to Nursing and Health Data Research, Arctic Ice Simulation from 1986–1990, and Experience with Autotasking a Seismic Inversion Code.

GLOBAL OCEAN MODELING ON MASSIVELY PARALLEL MACHINES

A high-resolution global ocean circulation model has been developed for the massively parallel Connection Machines (CM-200 and CM-5) at the High-Performance Computing Research Center by Los Alamos scientists Richard Smith, John Dukowicz, and Robert Malone.

Computer models that simulate numerically the behavior of the atmosphere and oceans are the best means we have for projecting future climate and the impact on it of human activities. The ocean is a critical component of the climate system. It plays a major role in the transport and storage of heat; and, because of its tremendous thermal capacity, it modulates the high-frequency variations of atmospheric weather. Changes in the patterns of ocean circulation may be responsible for transitions between different climatic regimes.

Development of a new generation of climate models that run on massively parallel computers is a major objective of DOE's CHAMMP and HPCC Programs. The increased computation power of massively parallel computers will make it possible for future models to have finer spatial resolution and more realistic treatment of the physical processes that control our climate.

The Los Alamos model, based on the widely used Bryan-Cox-Semtner ocean model, has been completely rewritten and extensively reformulated for efficient execution on the Connection Machine. It simulates the evolution in time of the ocean currents and distributions of temperature and salinity. The model domain is the three-dimensional global ocean, including realistic bottom topography and coastal boundaries of continents and islands. The model is forced by atmospheric wind stress and fluxes of heat and fresh water at the ocean surface; these can be obtained either from a global atmosphere model or from climatological observations.

Significant improvements have been made in the formulation and solution techniques of the parallel version of the model. Some advantages of these improvements are as follows. (1) Any number of islands can be included at no extra cost. For example, the new model uses all 80 islands that can be seen at the resolution of the picture, in contrast with previous calculations at the same resolution using the traditional method that was limited to only three "islands" (Australia, New Zealand, and Anarctica) to save computer time. (2) The old method requires that the bottom topography be smoothed substantially; the new method runs with no smoothing. (3) A new solution technique for massively parallel computers has been developed that greatly reduces the time spent solving the model. (4) The traditional condition at the ocean surface has been replaced with a more physical free-surface treatment that has the added advantage of even faster computation. In the picture, one simulated year takes about 23 hours on a full CM-200.

This work is sponsored by the DOE CHAMMP and HPCC Programs and is being carried out in close collaboration with climate scientists sponsored by the National Science Foundation and the National Oceanic and Atmospheric Administration.

Massively parallel computers run climate models that simulate the behavior of the atmosphere and the oceans. These models are the best means we have for projecting future climate and the impact of human activities.

PORE-SCALE MULTIPHASE FLOW IN COLLABORATION WITH MOBIL OIL

Nearly all US domestic petroleum reservoirs have already undergone primary production. Any subsequent production of domestic petroleum reserves requires more expensive secondary and tertiary production methods. The economic decisions to develop a reservoir are based in part on estimates of resources in the field and simulation of the extraction process using reservoir models. In the face of relatively inexpensive foreign oil supplies, domestic producers must have confidence in simulation results before risking investment in improved oil recovery methods. The field simulation codes used for this purpose are very sensitive to the interactions between the fluids and the rock. These interactions are not always well represented in the simulations. This is one significant reason that the record of improved oil recovery simulation has often been questionable.

We are simulating the pore scale hydrodynamics of multiphase flow to provide fundamental representations of the oil displacement processes. In this simulation, we are using data from real rock samples supplied by a collaborating petroleum research laboratory (Mobil Exploration and Production Technology and Engineering Center) to describe mathematically the interactions between these fluids and rock for use in reservoir scale models. We will then test these simulations by comparing them with laboratory and field-scale results also supplied by industry.

Lattice gas and lattice Boltzmann algorithms have been developed at Los Alamos and applied by the authors during the past two years to simulate numerous pore scale flow phenomena. This figure shows a computer simulation of a two-immiscible component, three-dimensional displacement of oil by water from the pores of an oil-bearing sandstone. The grains of the rock were obtained from x-ray CT scan images; the cube of rock shown is one one-hundredth of the complete data set. We expect to be able to solve the entire problem on the Thinking Machines CM-5 in the DOE HPCRC.

The interactions between fluids and rock can be described mathematically and simulated on a computer. The photos show displacement of oil by water from the pores of an oil-bearing sandstone.

PVM: HETEROGENEOUS DISTRIBUTED COMPUTING RESEARCH

PVM (Parallel Virtual Machine) is a software package developed at Oak Ridge National Laboratory that permits a heterogeneous collection of UNIX computers hooked together by a network to be used as a single large parallel computer. Thus, large computational problems can be solved by using the aggregate power and memory of many computers.

PVM makes a user-defined collection of serial, parallel, and vector computers appear as one large distributed-memory computer. PVM supplies the functions to automatically start up tasks and allows the computers to communicate and synchronize with each other.

Applications, which can be written in FORTRAN or C, can be parallelized by using simple message-passing constructs common to most distributed-memory computers. By sending and receiving messages, subtasks may run on a vector supercomputer and other subtasks may run on a parallel computer or powerful work station. At the machine level, PVM automatically handles all message conversion that may be required if two computers use different data representations. At the network level, PVM allows the computers being used to be distributed anywhere in the world and tied together by a variety of networks.

The PVM source code and user's guide are available by electronic mail. The software is easy to install and can be done by any user. The source has been tested on Sun, DEC, IBM, HP, SGI, and Next work stations, as well as parallel computers by Sequent, Alliant, Intel, Thinking Machines, BBN, Cray, Convex, IBM, and KSR.

PVM is an enabling technology. Hundreds of sites around the world are already using PVM to solve important scientific, industrial, and medical problems. DOE research laboratories using PVM include ORNL, LANL, LBL, ANL, PNL, Ames, INEL, as well as researchers at NERSC and NOAA. NSF Supercomputer Centers at Illinois, San Diego, Florida State, Utah, Pittsburgh, North Carolina, and Cornell all have many PVM users. NASA Research Centers using PVM include Langley, Ames, and Marshall. Numerous universities around the country are using PVM both for research and as a teaching tool.

In the areas of technology transfer, Cray Research, Convex, and IBM have decided to supply and support PVM software on their systems.

To receive this software, send e-mail to netlib@ornl.gov with the message, "send index from pvm." An automailer mail handler will return a list of available files and further

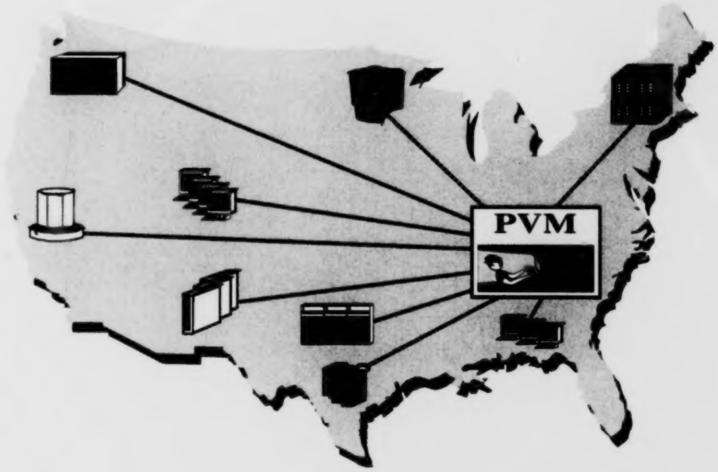

Computational Grand Challenges such as weather modeling use PVM (Parallel Virtual Machine) to exploit the aggregate power of heterogeneous workstations and supercomputers distributed around the nation.

instructions by e-mail. PVM problems or questions can be sent to pvm@msr.epm.ornl.gov for a quick and friendly reply.

ADVENTURES IN SUPERCOMPUTING

Adventures in Supercomputing (AiS) is the education component of ORNL's DOE-funded High-Performance Computing Research Center. The program is aimed at attracting women, minorities, and other under-represented groups to careers in science and technical fields. This high-school-level effort is being coordinated by ORNL, Ames Laboratory, and Sandia National Laboratories at Albuquerque.

The Tennessee AiS Review Board selected seven schools from among 39 qualifying applications from the state. The curriculum was taught to two teachers from each selected school during a Summer Institute in Oak Ridge. For a student computing environment, each school has received, on loan from DOE, four color Macintosh computers with visualization software and a color printer. Each school has Internet access through SURAnet to an nCUBE parallel computer at ORNL, which was provided by the nCUBE Corporation.

An assessment program has been developed to allow short-term and long-term monitoring of the progress of AiS teachers and students. The first phases of assessment were administered during the Summer Institute. Discussions have begun with Colleges of Education to help fill the pipeline with prepared teachers, and with the Tennessee Department of Education with respect to inclusion of computational science in the state curriculum.

P-FEM APPLIED TO A REAL FIELD SITE

Modeling the movement of hazardous waste in groundwater has been identified by the Department of Energy as one of the "Grand Challenges" in scientific computation that needs to be addressed in the next decade. In recognition of this need, DOE has provided support for a group of scientists located at several national laboratories and universities to conduct research and development in groundwater flow and contaminant transport modeling. This group is part of a larger consortium of researchers, collectively referred to as the "Partnership in Computational Science" (PICS) that has been charged with the task of applying high-performance computational tools and techniques to Grand Challenge areas identified by DOE. The results described here are the first application of parallel computing to this important problem area.

The site selected for P-FEM application is Melton Valley on the Oak Ridge Reservation, which includes Waste Area Group 6 (WAG 6), the only operational low-level waste disposal facility at ORNL. Because of concern over potential off-site migration of hazardous wastes, WAG 6 and other waste area groupings located in Melton Valley have been the subject of extensive site characterization and environmental monitoring activities. Calibration of the groundwater flow model is based on the substantial amount of data that has been collected from regular monitoring of piezometer and water quality wells installed within and around the ORNL site.

The conceptual model of the flow system consists of three layers: regolith, shallow bedrock, and deep bedrock. This conceptual model is based on results of geologic and hydrologic investigations. The finite element model reflects the identified geologic and hydrologic investigations. The finite element model reflects the identified geologic layers and serves as the computational grid for modeling of the Melton Valley site.

A model has been calibrated to field data (i.e., water level measurements) and is being used to address pertinent issues ranging from fundamental understanding of the Melton Valley groundwater flow system and identification of data needs (such as locations of additional wells) to evaluation of proposed corrective measures for remediating waste sites. This model was used to predict the effect of capping

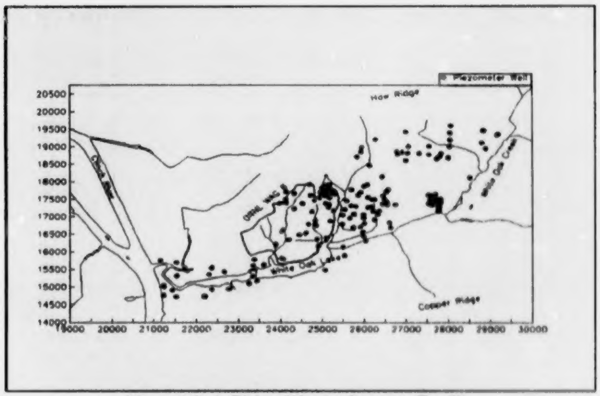
Piezometer Well
Haw Ridge
Clinch River
ORNL WAG
White Oak Lake
White Oak Creek
Copper Ridge
20500
20000
19500
19000
18500
18000
17500
17000
16500
16000
15500
15000
14500
14000
19000
20000
21000
22000
23000
24000
25000
26000
27000
28000
29000
30000

Map of Oak Ridge Waste Storage Facility

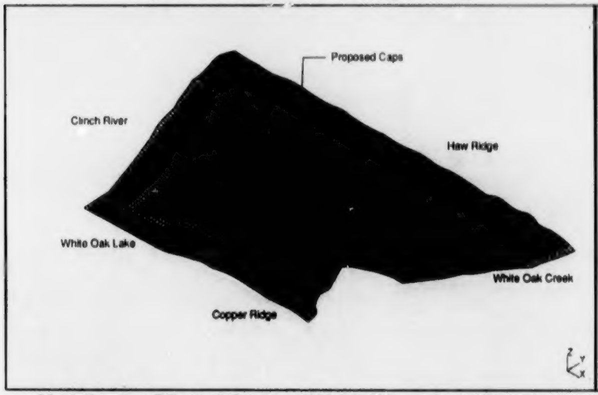
Proposed Caps
Clinch River
Haw Ridge
White Oak Lake
White Oak Creek
Copper Ridge
z
y
x

Modeling the Effect of Capping WAG 6 Waste Burial Trenches

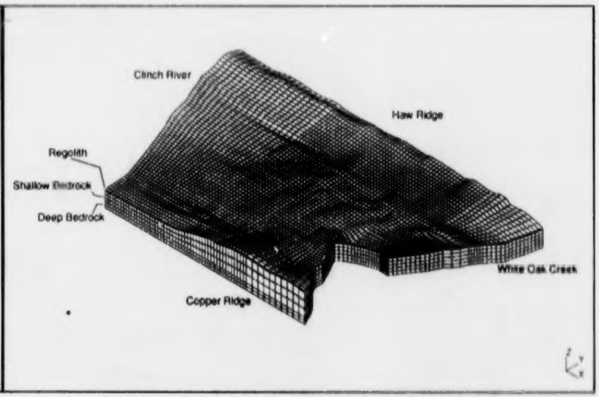
Clinch River
Haw Ridge
Regolith
Shallow Bedrock
Deep Bedrock
White Oak Creek
Copper Ridge
z
y
x

Computational Grid for Waste Area Grouping 6

LOAD SET: 1 TIMESTEP: 1 TIME: 0.0
Mesh2l: d1caps, recharge=2.4e-8 fps
WATER TABLE PROFILE MIN: 711.80 MAX: 1041.00
1041.00
860.00
840.00
820.00
790.00
780.00
760.00
750.00
740.00
711.80
z
y
x

Computed Water Table Elevations after Capping

P-FEM: Parallel 3DFEMWATER applied to a real field site.

WAG 6 waste burial trenches on the groundwater flow system. Proponents of this approach to remediating WAG 6 claim that caps will lower the water table and, in effect, will prevent hazardous wastes in burial trenches from coming into contact with the groundwater.

The model is being used to test the effectiveness of capping in lowering the water table. The figure displays the water table elevations computed from the model on the Intel distributed-memory parallel computer. Model results were used to identify WAG 6 wells that will be affected by capping. They will also be used to provide technical guidance for drafting a plan for WAG 6 Remedial Action Groundwater Monitoring Activities.

Further work will include modeling anistrophy (preferred flow direction) and how it would affect the flow system and visualization of three-dimensional flow fields to possibly uncover flow/contaminant pathways that have not been considered in the past (e.g., underflow beneath White Oak Lake). Flow fields will provide input to future contaminant transport modeling (modeling of tracer tests).

This model would not have been feasible without the high speed and memory capacity provided by the parallel machine. Running this problem on "slower" computers would have required much more time (several orders of magnitude more). Speed and quick turnaround time are particularly important for the calibration process. The speed provided by parallel machines would also allow quicker responses to queries by environmental agencies (EPA) regarding contaminant migration from DOE waste sites.

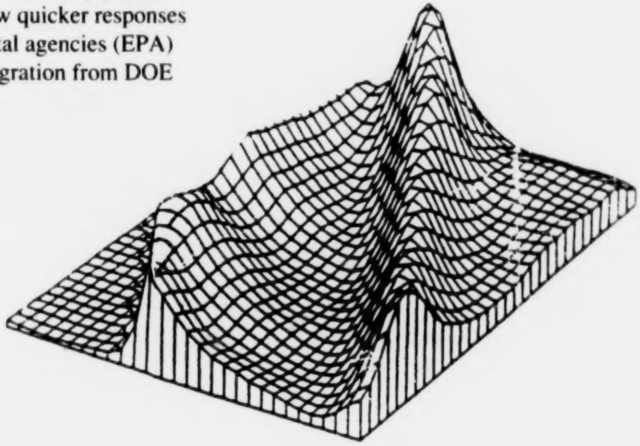

Sandia National Laboratories

Adventures in Supercomputing

Adventures in Supercomputing (AiS) is a collaborative program between Oak Ridge National Laboratory, Ames Laboratory, and Sandia National Laboratories. The program is aimed at attracting women, minorities, and other underrepresented groups to careers in science and technical fields.

The AiS New Mexico Summer Institute was held June 8–19, 1992, at Sandia's Massively Parallel Computing Research Laboratory. Seventeen math, science, and computer science teachers from eight high schools attended the two-week training session, at which they were introduced to high-performance computing and several areas of computational science. The teachers gained a working knowledge of the UNIX operating system, the FORTRAN programming language, the national research and education network, and available software, and were provided lesson plans, curriculum materials and course outlines for the upcoming school year.

A modified version of the classroom configuration for the AiS schools was used at the Institute. Each had four color Macintosh IIsi's with network connections and one Deskwriter color printer. Once back at the schools, the Macintosh local area network was linked to the national research and education network. Sandia worked with New Mexico Technet and the University of New Mexico to provide the networking hardware and access to the schools.

In addition to the Summer Institutes at the three participating Labs, AiS is sponsoring two Summer Faculty Fellows at Sandia this year. Dale Harris and Pam Lisle are working on computational science projects for use in the classroom and are collecting information about national computing and networking resources to share with their AiS colleagues in the fall.

New Mexico schools participating in DOE's 1992 Adventures in Supercomputing program were: Alamogordo High School, Albuquerque High School, Estancia High School, Gadsden High School, Las Cruces High School, Moriarity High School, Onate High School, and Vaughn Municipal.

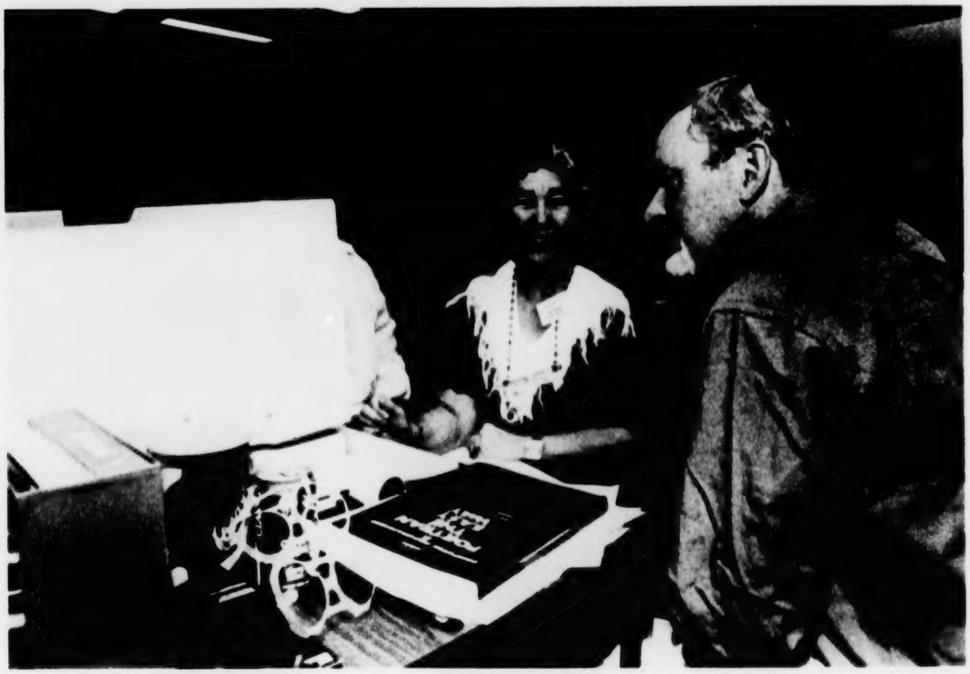

Sandia President Al Narath with Estancia High School teacher Mildred Pritchett.

Computer Design of Materials and Molecules

The goal of CDMM is to develop breakthrough capabilities in materials simulation using innovative new MP algorithms and methods and MP computers. The fundamental significance and broad economic, societal, and scientific impacts of this project will be realized through a strong interdisciplinary approach involving partners with expertise in catalysis, chemistry, and materials science as well as high-performance computing, and by coupling with experimental programs to validate the progress of computational results.

The partners are focusing on chemical and biocatalysis-catalysis for energy production, pollution minimization, and biotechnology. Some of our complementary efforts include micro- and nanoelectronic and photonic materials for computers and consumer electronics and novel structural materials for manufacturing. The computational techniques for these materials are closely related, thus allowing a coordinated attack on MP methods at three levels: first, development of scalable, microscopic quantum mechanical models (including Hartree-Fock, Local Density Approximation, and Many-Body methods); second, atomistic molecular dynamics and Monte Carlo methods for classical, microscopic force-field simulations; and, third, macroscopic continuum models.

Our five-year goal is to routinely produce fully detailed 3-D simulations for the continuum models, microscopic force-field calculations with tens of millions of atoms, and full quantum mechanical relaxations for systems with hundreds and thousands of atoms. A nested capability is planned in which the continuum models provide boundary conditions for the force-field calculations, and the force-field calculations are used to provide the boundary conditions for the quantum mechanical simulations. The full achievement of the five-year goal is predicated on the availability of teraflop and terabyte computers.

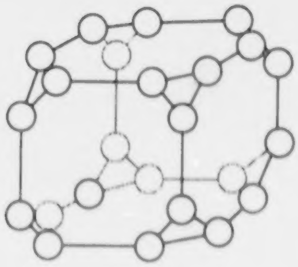

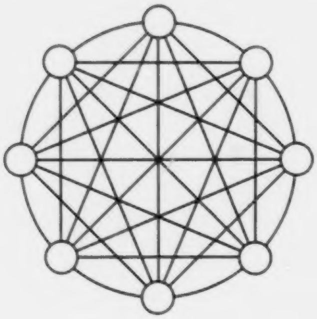

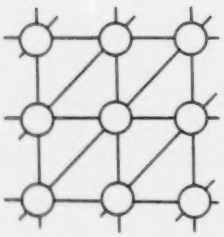

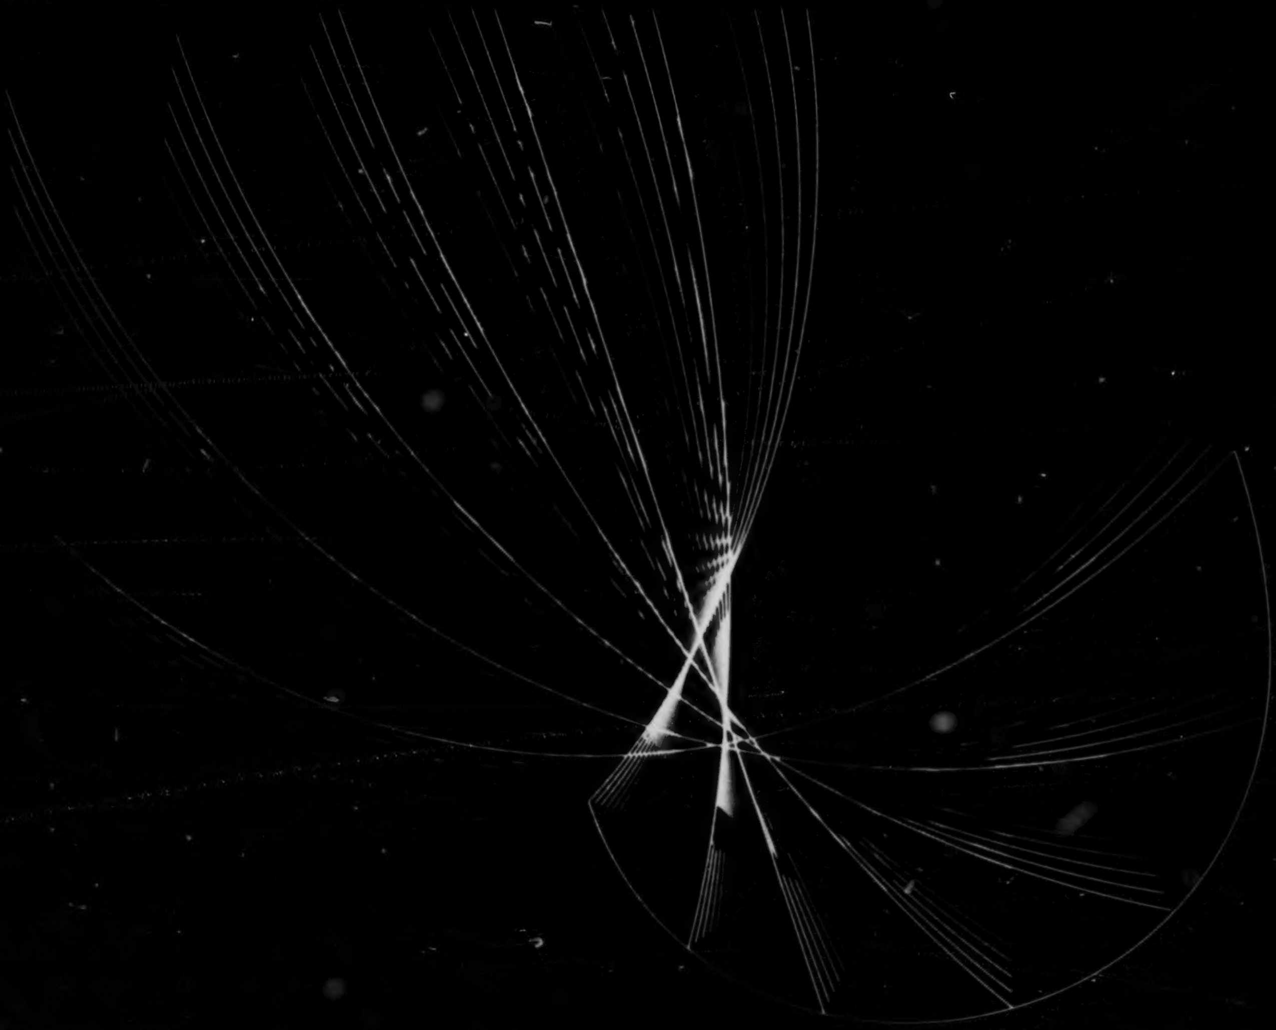

Section 3

Program Summaries

BEST COPY AVAILABLE

HIGH-PERFORMANCE COMPUTING SYSTEMS

GRAND CHALLENGES

The Grand Challenge effort at Ames is being carried out in the context of the Partnership in Computational Science centered at Oak Ridge National Laboratory. The effort at Ames is focusing on the development and application of molecular dynamics simulations for materials undergoing martensitic phase transformations. These are transformations between one solid crystalline phase and another, and they are very important in the hardening of steels and for titanium alloys used in aircraft materials. Significant progress has been made in studying Zr metal using the nCube parallel computer. This initial period has also involved finding personnel with strong computer science backgrounds to start converting algorithms to parallel architectures. A brief summary of the scientific activities is given below.

Our experience with total energy calculations suggested that the interatomic forces in Zr metal might be adequately modeled with a classical potential. The martensitic transformation in this system has been extensively studied experimentally and there is great interest in knowing if the dynamical behavior of the atoms near the phase transformation somehow signals if the lattice is about to become unstable. We adopted the embedded atom method and ran molecular dynamics simulations on several serial computers to establish benchmarks. The code was then modified (over many months) to run on the nCube. It was successful in making runs with up to 72,000 Zr atoms. For the dynamical information of interest, runs with 3,000 to 10,000 atoms were adequate; but the larger number of atoms will be important for answering other important questions of interest to metallurgists (e.g., the influence of dislocations). The time required scaled linearly with the number of atoms and inversely with the number of processors. There was some loss of efficiency with 256 processors, but saturation had not yet occurred. The speed of 32 nCube processors was equivalent to 1 processor on the CRAY 2.

Considerable effort has been spent exploring schemes to include the electronic degrees for freedom into the molecular dynamics (i.e., going beyond the classical interatomic potentials). One possible method is to use a parameterized tight-binding Hamiltonian, and we have developed a tight-binding molecular dynamics scheme specifically for simulating systems with directional bonding. To make the algorithms scalable, we have to go beyond the traditional diagonalization methods and use instead a Green's function formalism. The first attempt was to employ the recursion method. Initial tests indicate that the method converges rather slowly with the number of atomic shells included, and thus it may not be economical for simulations that require a high degree of energy conservation over 10,000 time steps. In the next stage, the scalability of other algorithms such as spectral techniques and the resolvent matrix method will be tested. Initial steps were also taken to consider ab initio simulations using the local density functional formalism.

PERFORMANCE ANALYSIS

Researchers work in several areas of computer performance analysis: fixed time benchmarking (SLALOM), multiparameter performance models for understanding and prediction of an application-algorithm-architecture triple, dynamic load balancing, provably optimal dataflow graphs via exhaustive search, and mathematically sound definitions of computational "work." The SLALOM work has advanced greatly in its realism and the level of understanding it imparts; it serves as our best-understood application program when dealing with new programming models and tools. The performance models are part of any computational science papers we write; we do not regard it as an add-on area of study. Finally, provably optimal algorithms and rigorous definitions of computational "work" represent the most long-term and difficult research area in the group. Casual efforts in the past have not succeeded in making computer performance analysis into a science; instead they have littered the field with ill-thought-out metrics. The guiding philosophy of the effort in performance analysis is to attempt the same level of scientific method practiced by competent experimentalists in the physical sciences.

COMPUTER SYSTEMS PERFORMANCE ANALYSIS PROJECT

This activity is designed to provide a coherent computer systems performance analysis program with the objective of improving the design and utilization of next generation computer systems for scientific and engineering computation. Program elements include the following activities.

1. An Industrial Computing Requirements Forum to improve communication between systems designers and industrial (and other) users.
2. A Grand Challenge Performance Analysis effort to establish and utilize better methods for analyzing the performance of GC applications codes.
3. A program in Applications Education for Computer Engineers to assist in better educating computer engineers about the scientific and engineering applications of computer systems.
4. An activity in Software Design for Advance Architectures that is designed to provide curricular tools to assist in teaching people correct programming practices for use on parallel computers.

EVALUATION OF EARLY SYSTEMS: MASSIVELY PARALLEL METHODS FOR ELECTRONIC STRUCTURE CALCULATIONS

We are exploring techniques for the solution of Schroedinger systems on massively parallel architectures, with the objective of studying cohesive and structural properties of metals, alloys, and intermetallic compounds. In particular, we focus on the use of optimization methods with preconditioners and on the development of scalable, multidimensional fast Fourier transforms. We are working with researchers in the Materials Science Division at Argonne, at Lawrence Livermore National Laboratory, and at the University of Illinois, as well as with industrial collaborators at Thinking Machines Corp. and Universal Energy Systems Inc. To date, we have developed a one-dimensional pseudo-spectral code for the Connection Machine; conducted a study of performance of one-, two-, and three-dimensional complex-to-complex FFTs and several real-to-complex/complex-to-real routines (for use in pseudo-spectral methods); and explored a Hamiltonian approach used in electronic structure calculations.

We also continue to operate the Advanced Computing Research Facility, through which Argonne researchers have access to the Intel Touchstone DELTA. The ACRF machines also directly support research on advanced software technology and algorithms, CHAMMP climate modeling studies, and research for the NSF Science and Technology Center for Research in Parallel Computation.

Los Alamos National Laboratory

ADVANCED PROTOTYPE SYSTEMS

The High-Performance Data System (HPDS) Project addresses the need for fast, reliable, high-capacity data storage and retrieval. Advances in massively parallel, large-memory computers and cooperative processing networks allow researchers to execute large-scale codes that generate massive amounts of data. Recent developments in high-speed networks at Los Alamos National Laboratory (LANL) allow large amounts of data to be transferred through a HIPPI (High-Performance Parallel Interface) network at high speeds. The HPDS, currently in development at Los Alamos, will consist of a collection of high-performance storage systems connected directly to a HIPPI network and managed by workstations. Data will be transferred directly between the storage systems and client computers instead of the traditional method requiring an intermediary mainframe computer. New technology disk arrays and new technology tape systems will be used to achieve the HIPPI transfer speeds and the capacity necessary to meet user requirements.

A prototype disk storage system is in development and an early test version has been used as "proof of concept" for some of the basic ideas in HPDS and for demonstration of high-speed data transmission (up to 61 Mbyte/s transferring images to a frame buffer) using HIPPI-attached storage devices. Data compression techniques are being developed that will reduce the data storage and transmission requirements. Based on the current rate of growth of storage technology (disk and tape), a factor of 10 compression could lead to a five-year gain on storage technology enabling larger computations on a shorter time scale.

Evaluation of Early Systems

As charter members of the Cray Research MPP design group, LANL staff have had a marked influence on the design of a new generation of massively parallel machines (the Cray T3D). As we have found, current-generation massively parallel machines are seriously limited in the types of problems that they can efficiently compute. The goal of the Cray T3D Early Evaluation Project is to investigate and resolve these limitations in the context of the Cray T3D system. The T3D has a number of unique components that distinguish it from current-generation machines. This joint effort between LANL and Cray will continue to investigate the strategies for the use of these components and evaluate their usefulness to the performance of important DOE applications. LANL brings unique qualifications in large-scale applications development, supercomputer performance evaluation, and active membership in the MPP design team to this proposal.

ADVANCED DATABASE ARCHITECTURES

We are pursuing two projects under the HPCC Program. Both are concerned with managing massive amounts of data on tertiary storage. Current database management systems (both relational and object-oriented) interface only to disks. Thus, the management of tertiary storage through a standard interface will be beneficial to both projects. One project (in conjunction with Argonne National Laboratory and the University of Illinois) already was started last year. Its purpose is to develop efficient access to high-energy "event" data (data that results from collider experiments). This project's goal is to organize the massive amount of data expected from SSCL experiments in order to facilitate efficient analysis. The second project, started this year in conjunction with Lawrence Livermore National Laboratory, aims to develop efficient access to massive amounts of spatial and temporal data. Initially, it will concentrate on climate modeling data.

The above two projects share the goal of organizing the data in such a way as to optimize its access for subsequent analysis. The main idea is that the data sets will be partitioned as to reflect their most likely access. For example, spatio-temporal data can be organized according to temporal partitions if this is the most prevalent access pattern. However, the nature of the data is so diverse in these projects that different techniques are necessary for their partitioning into smaller subsets. In the case of physics event data, it has been demonstrated that organizing the data according to their properties over all events resulted in an order of magnitude improvement in access time. In the case of climate modeling data, we need to develop algorithms for partitioning and re-assembly of the data along multi-dimensional boundaries. Another difference is that in the case of the climate modeling application, there is existing software that requires interface to files. Thus, the application software will interface directly to the specialized software that partitions and assembles the datasets. In the case of the event data, it will be managed by a database management system (both relational and object-oriented systems are being experimented with), and thus a programming language interface to the database system will be used.

These two projects are expected to share technology. The partitioning of event data is described by using an object-model (EER), and thus our database tools are already being used for the relational implementation of the object structure. Similarly, we expect to use the object-level tools to describe the partitioning information as well as the metadata for the climate datasets. This data will be managed by a relational system and used by the partitioning and re-assembly modules. Another important shared benefit is expected from the "mass storage server" that will be developed in conjunction with Lawrence Livermore National Laboratory. Our plans are to extend the current protocol of the Unitree server so as to have control over the placement of individual subsets on tertiary storage. We expect this technology to be useful to the physics datasets as well, since the enhanced Unitree could be interfaced with externally to the database system that manages the event data.

EVALUATION OF EARLY SYSTEMS

The purpose of this project is to study innovative hardware and software features of new computers for Grand Challenge codes and to provide feedback to machine architects and system software designers. Two major subtasks comprise the project: benchmarking and predictive modeling.

The goal of the benchmarking task is to explore achievable computation rates and what is required to achieve these rates. In FY92, our major effort involved testing and evaluating serial number one of the Kendall Square multiprocessor. In addition, we tested and evaluated message-passing on the Intel DELTA mesh, analyzed the communications performance of the Intel iWarp, and began the testing and evaluation of an Intel Paragon beta unit.

The primary goal of the predictive modeling task is to construct predictive models of parallel systems. In FY92, we focused upon two architectures, the Intel iPSC systems and the Kendall Square KSR-1. Experimentation, simulation, and analytical modeling activities were performed on various Intel systems with the motivation of determining the optimal number of processors to allocate to an application, characterizing workloads, and automatically identifying distinct phases within application codes. Modeling and experimentation are also underway on the KSR-1 system, focusing initially upon the "post store" memory option, which can have dramatic effects on the performance of the overall system. Analytic models are being constructed and validated.

Sandia National Laboratories

HIGH-PERFORMANCE COMPUTING AT THE MPCRL

Sandia researchers have increased and broadened the impact of high-performance computing by applying results of our research programs in mathematical algorithms and computer science interdisciplinary computational research in chemistry, materials science, engineering, and physics. Some of the largest benchmark quantum mechanics and molecular dynamics computations conducted to date were performed on Sandia's massively parallel (MP) computers. A collaboration with BIOSYM developed an improved understanding of dihydrofolate redutase enzyme; this work has potential impact in biomimetic catalysis, basic science (especially biochemistry), bio-technology and medicine (cancer chemotherapy). Sandia received a national award from the microscopy community for developing revolutionary new capabilities for electron trajectory simulations. In collaborations with NASA and Boeing high performance computing technology was used to analyze and help design bumper shields that protect space vehicles from collisions with space debris. MP computing provided critical simulations for Operation Desert Storm; these simulations could not have been conducted in a meaningful time frame without Sandia's models, MP algorithms, and MP computers.

MP eigenvalue solvers and FFTs developed at Sandia are used in large quantum mechanical simulations. Highly efficient, scalable iterative solution algorithms for sparse linear systems are used to solve high-Mach-number flow problems and will form the basis for new combustion models. New dynamic load balancing and heterogeneous programming strategies made possible the Desert Storm computations. These load balancing algorithms were also applied to electron trajectory simulations and the heterogeneous programming strategies will, in the future, allow researchers to perform coupled atmosphere and ocean simulations on a single MIMD computer. Advanced domain decomposition methods made possible the successful implementation of unstructured finite-element methods on MP computers and are being applied to problems in combustion and ocean simulations.

A parallel UNIX-based operating system developed at Sandia is providing increased functionality for applications and has enhanced portability between different MIMD machines. A new graphical performance monitoring tool is complementing efforts at other laboratories by providing, for the first time, real-time display and scalability to thousands of processors. Researchers received national attention for work in network security, network interoperability and gigabit networks. Sandia integrated diverse technologies to produce a 19-kilometer high speed (1.2 gigabits per second) network link during Supercomputing '91; applications executed remotely on a CM-2x at Sandia's Massively Parallel Computing Research Laboratory as visitors watched the results at the Albuquerque Convention Center.

Education is an important thrust at Sandia. Our K–12 activities include Adventures in Supercomputing, SuperQuest, New Mexico High School Supercomputing Challenge, and the first Mathematics Through Applications Workshop. A number of Sandia staff spend a day a week working in local schools as science advisors and a summer research program for teachers has been initiated. Sandia staff participate in the FAME (Fellows for the Advancement of Mathematics Education) Program involving parents in the mathematics education of their children. The goal of these programs is to improve science and mathematics education through the aggressive use of computing.

ADVANCED SOFTWARE TECHNOLOGY AND ALGORITHMS

SOFTWARE TECHNIQUES AND ALGORITHMS

This program element concentrates on the issues of matching science and engineering applications to modern computer architectures and algorithms. Applications currently include quantum wave scattering, image synthesis, particle transport, semi-empirical computational chemistry, n-body simulations, computational fluid dynamics, and mesoscale climate simulation. The architectures we apply are scalable, distributed memory parallel computers that we feel represent the trend of high-speed computing: MasPar, nCUBE, and Intel are the on site equipment base. The focus on algorithms is to find very novel and aggressive methods that give several orders of magnitude performance improvement instead of simply porting traditional algorithms to run in parallel. The combination of parallel computing and improved algorithms is proving very powerful for the aforementioned applications, all of which have defied conventional computational approaches.

UNDERGRADUATE EDUCATION PROJECT

The primary purpose of this project is to develop textbooks that introduce computational science at the middle school level (6th–8th grades) and at the level of a high school senior or college freshman. The latter will be a coordinated activity that parallels the graduate level text on computational science that is being supported by the SCF. The texts in question will have as a component module software that in part may be based on projects that evolve from the Adventures in Supercomputing component of the DOE HPCCP. The purpose of these modules is to extend and illustrate the traditional textbook material.

ACCESS TOOLS

This project develops software tools that assist in using high performance computers for grand challenge applications in a distributed environment. Online tutorials that allow researchers access to new machines, tools that provide access to hardware and software information, and pointers to individuals working within Grand Challenge applications are being developed. The goal of this project is to replace simple linear text, e.g., digital books, with information resources that more fully realize the potential of computing environments for information access and transfer.

Argonne National Laboratory

SOFTWARE SUPPORT FOR GRAND CHALLENGES

Argonne is conducting several projects to enable the use of high-performance computer architectures to solve Grand Challenge problems. In computational chemistry, we are developing and testing new algorithms and modeling systems that scale to larger numbers of processors. In particular, we are quantifying tradeoffs between accuracy, parallel scalability, absolute performance, and memory requirements. This information—which is obtained by a combination of algorithmic analysis, mathematical modeling, and prototyping—will enable us to design chemical modeling systems with a good degree of confidence that they will remain competitive on future generations of parallel supercomputer. Our work is driven by the requirements of problems in the chemistry of halohydrocarbons, the chemistry of clay minerals, and the rational redesign of biodegradative enzymes.

In computational biophysics, we are developing new mathematical models and algorithms to study protein-protein interactions. Our primary effort is directed to the development of a parallelized version of XPLOR and AMBER, the principal software tools of protein scientists. We are also exploring the potential benefits of advanced scientific visualization techniques, in particular, volumetric display and virtual reality systems. Our goal is to develop the technology to accurately model the interactions of two proteins and to predict the changes in interaction that occur as the result of engineered substitutions of amino acids involved in the interaction. This goal is the first step to create the technology needed for a true molecular engineering capability. In materials sciences we are focusing on vortex dynamics in type-II superconductors. We have developed a model of a single flux vortex as a filament under tension, subject to impurity potentials and thermal fluctuations. This model gives good quantitative insight into vortex behavior in the presence of disorder for a wide range of temperatures and magnetic fields and provides a first opportunity to investigate glassy dynamics under small applied forces and scaling behavior near dynamic phase transitions. We have also been investigating the dynamics of vortices according to the time-dependent Ginzburg-Landau equation. In this effort, we have successfully parallelized a three-dimensional program to run on the Intel DELTA using a domain decomposition approach.

SOFTWARE COMPONENTS AND TOOLS

Supporting these efforts is research in software tools. We have launched a four-pronged research program based on programming parallel computers, reusing parallel code, developing scalable algorithms, and understanding performance. The basis for these efforts are the performance-monitoring tools and PCN. For example, we are using the PCN system to develop rapid prototypes for promising algorithmic alternatives in the computational chemistry project. We are also using pp to parallelize the CONGEN protein modeling program.

In another area of tools development, we are developing automatic differentiation techniques that allow one to compute derivatives of large codes accurately and efficiently. Our software package ADIFOR is currently being used by engineers in engine and airfoil design, and our package ADOL-C is being incorporated into an advanced simulator for the simulation of complex vehicles in real time.

COMPUTATIONAL TECHNIQUES

We are developing multidimensional fast Fourier transforms for the study of ternary intercalation compounds. These compounds include transition elements with from 50 to 100 atoms. Although such problems are in the range of those treated by conventional supercomputers, the high cutoff energies significantly increase the computational problem.

FLUID FLOW IN POROUS MEDIA: CONTAMINANT TRANSPORT IN GROUNDWATER

This portion of the program is part of a collaborative effort, the PICS consortium, funded under the Advanced Software Technology and Algorithms component of the High Performance Computing and Communications Initiative. Together with groups at ORNL, Stony Brook, Rice University, The University of South Carolina, and Texas A&M, a multi-year program is under way to develop significantly improved, fully three-dimensional models of groundwater transport which can be used to improve environmental remediation efforts. The efforts of the BNL group are focused towards the development of more robust and efficient elliptic solvers implemented for distributed memory parallel architectures. These new computational techniques will then be incorporated into the sequence of progressively more capable codes being developed under the overall program. Eventual application of these state-of-the art codes to real problems is an important goal; the BNL group has a close association with a joint interagency (DOE, EPA, NRC) effort in assessing currently used models, their shortcomings, and benchmark data for testing future models.

FIRST PRINCIPLE CALCULATIONS OF MATERIALS PROPERTIES

Funded also as part of the PICS consortium Grand Challenge Research program, this BNL program involves a close collaboration with the Solid State Theory group in BNL's Physics Department, and a somewhat looser collaboration with researchers at Oak Ridge and Ames Laboratories. The work involves the attempt to compute macroscopic properties of materials starting from a fundamental quantum mechanical description in terms of the ions and electron structures that are present. For these systems, the molecular dynamics is relatively straightforward: the challenge lies in obtaining accurate results for the electronic behavior. The BNL effort involves the use of the Car-Parrinello approach, which uses an iterative self-consistent effective potential approach. The goal is to extend previous work to allow these methods to be applied to systems of up to 1000 atoms, whereas present methods cannot deal effectively with more than a few tens of atoms. The approach is to sharpen the details of the numerical methods and to develop parallel codes. Specific objectives have been the development of well-behaved iterative diagonalization algorithms, which avoid the necessity of full matrix storage, which can be easily implemented as parallel codes, and whose accuracy and convergence properties are well understood.

Lawrence Berkeley Laboratory

SCIENTIFIC DATA MANAGEMENT

Various scientific projects at DOE need to support vast amounts of data generated by complex scientific projects such as environmental cleanup, climate modeling, human genome mapping and sequencing, and the superconducting supercollider (SSC). The amount of scientific data generated is accelerated because of more sophisticated devices and cheaper computer power. A single scientific experiment can generate hundreds of megabytes of data within days. Many scientific simulations are not carried out to the desired granularity level only because it is impractical to process and manage the large amounts of data that would be generated. The management and coordination of these large datasets pose significant challenges that cannot be readily met by existing commercial data management systems.

Over the last few years, the Data Management Group at LBL has been involved with several scientific projects that require data management support. These include databases for human genome, epidemiological, and superconducting magnet applications. Our research program is oriented toward the development of data management techniques and tools in support of such scientific applications.

Because of the special requirements of scientific applications (such as temporal, spatial, and sequence data) it is necessary to use database technology that can be extended to accommodate new data structures and operators. However, because of short-term needs and practical manpower considerations, we have adopted an evolutionary approach to the support of scientific applications. In the short term, we are developing techniques and tools that make it possible for scientists to use existing commercial relational database technology. At the same time, we need to have a way of migrating, in the future, to new extensible technologies. This is achieved by insulating the application programs from the specific database system used by using an intermediate object model. In the short term, we are developing translators from the object model to commercial relational data management systems. In the future, we plan to implement the same object model with new extensible database systems. This will provide a smooth migration environment since existing application programs will continue to interact with the same model.

Our work concentrated initially on the identification and characterization of the functional requirements of scientific applications. Based on these characterizations, we have developed conceptual models and physical data management techniques to support these requirements. For example, we have developed conceptual models for temporal data and multidimensional data, and various physical methods for the efficient implementation of these models. In general, once a conceptual model has been defined for a special scientific need, then methods for its efficient implementation need to be developed. Such methods can then be implemented as separate modules that can be used in conjunction with commercial relational systems to augment the object-level model.

Our current and planned work with several scientific applications is described below.

Database Tools

We are involved in two aspects of the Human Genome Project at this time. The first is a Chromosome Information System (CIS), which uses an object-level model to describe the application. The object structure (schema) is fairly complex and contains about 35 objects. We have used database tools that we have developed to translate this object structure into relations. The resulting relational schema was over 3000 lines of code in SQL and the trigger language of Sybase. The generation of such code by hand would have been an error-prone and complex task that would require a database expert. Instead, a biologist familiar with the application was able to describe the schema graphically and concisely in a matter of hours. We are currently developing query tools to access the object-level directly.

LIMS

Another activity associated with the recently started Human Genome Project is the Laboratory Information System (LIMS). In this project a user interface is being designed for describing a sequencing and/or mapping process in the laboratory. A translator will be developed to map the process flow diagram into equivalent object-level structure, which will be then supported by the object-level tools we have developed. This is an example of using the object-level translation technology to automatically support the application-specific model.

CEDR and SSCL

The above object-level tools have been used for a couple of other projects. One is the Comprehensive Epidemiological Data Resource (CEDR), whose goal is to collect all the information on low-level radiation in DOE in a single repository and to make it available for further analysis to researchers inside and outside of DOE. An interim system was developed using the current tools, and a future system will use the object-level query translator currently being developed. A second project using these tools is the SSCL Magnet Laboratory, mainly for developing the object-level schema and its translation to the Sybase relational database.

SUPERCOMPUTER ACCESS TOOLS

The goals of the Supercomputer Access Tools effort are to increase the research potential of high-performance computing. Team members work closely with scientists to identify areas where advanced computing technologies would be of greatest benefit. Collaborations include the LBL Magnet Measurement Engineering Group, the SSCL Magnet Test Laboratory, the LBL Human Genome Center, Nuclear Medicine, and the LBL Center for X-Ray Optics. Technologies currently under investigation include system integration tools, fault location techniques, high-speed interprocess communications mechanisms, icon-based programming paradigms, and network distributed shared memory as a potential interprocess communication mechanism.

DISTRIBUTED COMPUTING APPLICATIONS AND TESTBEDS

The Imaging and Distributed Computing Group of LBL currently consists of several staff scientists and students who research various aspects of computer imaging (including image acquisition, image storage, image processing, image understanding, image analysis, 3-D graphics, animation) and the use of high-speed distributed computing technology in imaging, storage systems, and data collection. In addition to its research the ITG assists LBL scientists in identifying imaging techniques that will address the imaging problems that arise in a wide variety of scientific disciplines. Some of the current applications include the following:

1. The ITG has developed several applications that demonstrate the use of high-speed networks to solve imaging problems. These applications have been used as instruments to explore the behavior of high-speed networks. The goal of all of these projects is to develop and promote the technologies required to build configurable distributed systems that will provide location-independent access to data collection, analysis, and storage facilities. For example:

1.1. One project uses distributed, heterogeneous systems to permit interactive manipulation of very large medical images. To create a visual representation from a typical 3-D magnetic resonance imaging (MRI) data set takes about five minutes on a typical scientific workstation. To be useful in a clinical or research environment, this speed needs to be reduced to less than one second. We have developed a prototype distributed system using a CRAY Y-MP, a CM-2, and an X-terminal connected via high-speed networks, which gives the required speeds. This work revealed and addressed many bandwidth limiting problems in several of the technologies involved in the interprocess communication. The work was also incorporated into a research demonstration presented at Supercomputing '91. We are currently working on other models of heterogeneous computing in order to determine and address the issues that limit high-speed distributed computing.

1.2. The ITG and ICSD are part of the DARPA-sponsored MAGIC Gigabit Network Testbed. MAGIC is undertaking the development of an application environment that emphasizes remote sensed imagery being delivered from distributed databases to visualization applications at data rates that permit real-time navigation.

The first test application will allow the simulation of walking, driving, or flying through a representation of landscape that is created from satellite and other images of the actual area of interest. The application is being developed in cooperation with the US Army's Future Battle Lab in order to test the hypothesis that field officers could benefit from this sort of capability in tactical situations.

The technologies needed to support this application are those of high-performance graphics workstations, high-speed distributed data storage systems, and wide area gigabit networks. The ITG is working on the various problems involved in designing and constructing a high-speed, network-distributed image server that will deliver images at the high data rates (400-600) needed to produce the visual representation of a real-time point of view moving through a complex landscape.

1.3. The ITG is collaborating with the Tenet network and RAID storage system research groups in the University of California, Berkeley Computer Science Department. This DOE-sponsored project will connect a digital video camera-based microscope to a gigabit network, and hence to the RAID-II network storage system. The network uses a combination of HIPPI local area and ATM wide area, network technology. An application is being developed to allow the researcher to capture and store the video stream in real time, and subsequently view and edit the video stream once directly from the RAID system. This scientific video data, unlike teleconference video, cannot be compressed using any of the common, glossy compression hardware since the compression inhibits subsequent image processing and analysis. This work is also part of a collaboration with the Blanca/XUNet gigabit testbed.

2. The ITG is working on automatic and human-assisted methods for segmenting various types of data, including MRI and videomicroscopy data. Significant progress is being made on software and algorithms that should greatly assist in the task of analyzing complex two- and three-dimensional image data sets.

3. The ITG is involved in programs to use sophisticated image-based applications to provide K-12 students with the involvement and motivation to learn computer skills and enhance the curriculum in a variety of disciplines, especially science and math. An example of this is the ITG "Whole Frog" Project. This project introduces the concepts of modern, computer-based 3-D visualization, and at the same time provides the power of whole body, 3-D imaging of anatomy as a curriculum tool.

The goal of the Whole Frog Project is to provide high school biology classes the ability to explore the anatomy of a frog by using data from high-resolution MRI imaging and from mechanical sectioning, together with 3-D surface and volume rendering software to visualize the anatomical structures of the intact animal. Ultimately we intend to be able to "enter the heart and fly down blood vessels, poking our head out at any point to see the structure of the surrounding anatomy."

The secondary goal of this project is to introduce the concepts of modeling and displaying 3-D structures directly from 3-D images obtained, for example, from MRI imaging (biological specimens), x-ray CT imaging (industrial imaging of non-biological objects), and direct generation from mathematical descriptions. This approach is being developed in conjunction with summer programs involving high school teachers and students.

Lawrence Livermore National Laboratory

COMPUTATIONAL FLUID DYNAMICS AND COMBUSTION DYNAMICS

This Grand Challenge project develops new numerical techniques for modeling combustion dynamics in realistic engineering configurations. This collaboration involves participants from the University of California, Berkeley; Lawrence Livermore National Laboratory; Courant Institute; New York University; and Los Alamos National Laboratory. The approach combines high-resolution finite difference methods with adaptive techniques that focus computational effort where required to minimize computational cost. The initial stages of this project will focus on three prototype applications that are representative of a broad range of combustion applications.

The first problem will be the modeling of a pulsed combustor in which acoustic effects play an important role. Pulsed combustors exhibit high efficiencies combined with low NOx emission, making them an excellent choice for heating systems. The second problem will be the simulation of a low NOx commercial burner. A medium scale prototype of this burner is under construction an Sandia National Laboratories, Livermore. This problem is a low mach number flow in which fuel (typically natural gas) and an oxidizer are ejected from a co-flowing nozzle to mix and burn inside a combustion chamber. The effect of nozzle design on turbulent mixing plays a key role in designing these types of systems. The third problem is that of detonations with realistic chemical kinetics. Understanding detonation behavior with realistic chemistry addresses an important safety problem in facilities that use combustible materials. In each of these problem areas, the goal is to develop computational capabilities suitable for engineering design and analysis.

NUMERICAL TOKAMAK PROJECT

Present day tokamaks, a particular design of a thermonuclear fusion energy source, are at the threshold of demonstrating the long sought goal of scientific "break-even"—producing as much energy as it takes to run the experiment. To develop these systems into successful commercial sources of electrical power is going to be a very costly process. Through this project high-performance computing will play a profoundly important role in designing these machines. The research consortium for this project involves scientists and engineers from six national laboratories, major contractors of the DOE and NASA, and five universities.

LLNL has begun research to create a predictive simulation of tokamak plasma turbulence in connection with the Numerical Tokamak Project. The work includes the creation of computer visualization diagnostics for the distributed interactive analysis of results. An important aim is to share scientific results over a high-speed, wide-area network to support collaborative work over long distances. Computer supported cooperative work on the Numerical Tokamak Project should stimulate new working arrangements among researchers, which will benefit other scientific and industrial projects as well.

THE SISAL LANGUAGE PROJECT

The Sisal Project focuses on the design, implementation, and use of a parallel language for large-scale scientific computing on next generation, massively parallel computer systems. This research is based on the premise that functional languages (such as Sisal) provide a highly cost-effective and portable method for writing parallel programs and substantially enhance the potential for compilers to exploit a diverse set of parallel machines.

The FY92 Sisal activities included work to improve the performance of compilers, enhance the interface to programmers, and expand the number of application users. The new, high-performance vector/parallel Sisal compiler outperforms automatic vectorizing and parallelizing FORTRAN compilers on Cray computers, Sequent Symmetry, SGI Iris, and other shared-memory machines. Sisal codes run as fast as hand-written parallel FORTRAN codes. Development has begun on Sisal systems for distributed memory computers. New heuristics based on "extended" dataflow graphs are being developed for program partitioning and task scheduling on NUMA architectures. Preliminary studies show this approach out-performs extant methods by 20% to 30%. The reduced cost of writing Sisal programs was further demonstrated when five versions of the 1-D FFT algorithm were developed in three weeks; the fastest code outperforms the parallel FFT library routine on the CRAY Y-MP.

The Sisal Scientific Computing Initiative (SSCI) was announced to expand the user community for Sisal. The initiative will promote the use of Sisal, educate the scientific community in the art of functional programming, create a self-sustaining user community, and begin the process of technology transfer. The initiative has grown to 35 participants, giving a total of over 70 sites using Sisal for application development work, research, or teaching.

OPTIMIZING MASS STORAGE ORGANIZATION AND ACCESS FOR VERY LARGE SCIENTIFIC DATASETS

This project will develop urgently needed scientific data management tools and integrate them into commercial mass storage and database management systems. The increased volume of data generated by modern large-scale supercomputer simulations and physical experiments overwhelm conventional methods of storage and access. The work focuses on the need of researchers from the climate modeling community. Initial efforts will support spatial and temporal climate modeling data. The project will develop data partitioning strategies based on the analysis of data access patterns and storage device characteristics, enhancements to current storage server interfaces and protocols to permit control over data placement on physical storage devices, use of customized data compression methods, and data re-assembly techniques for assembling the desired subset.

Contributions from collaborating institutions will be integrated into a prototype system using the National Storage Laboratory (NSL), housed at NERSC. Data management researchers at LBL will investigate data allocation and retrieval strategies. The University of Maryland will focus on data reduction and spatial organization. LLNL provides the operational and organizational environment for the proposal work, as well as the scientific applications and experimental datasets. Initial work has started at LLNL on "proof-of-principle" experiments that use pre-partitioned data clusters and predetermined access patterns to measure improvements in access speeds and to identify bottlenecks.

Work has also begun for design of new storage system functionality to place multiple clusters on a dynamic hierarchy of storage devices based on a desired service distribution. Changes will be made to several climate modeling applications to allow them to use the new storage system interfaces and protocols. The FY92 funds will be used to complete initial development of the storage system interface/protocol enhancements and climate modeling application changes. The integration of these enhancements and changes within the NSL will provide a large improvement in access efficiency for targeted climate modeling analysis and visualization applications.

SOFTWARE SUPPORT FOR GRAND CHALLENGES

Grand Challenge-scale applications are an essential ingredient of the DOE High-Performance Computing Research Center (HPCRC) at Los Alamos. Not only are they complex scientific problems of national importance, but also they are the primary driving forces for the development of the Grand Challenge Computational Environment (GRACCE). Some of the basics resources required for these simulations are given in the following table. Here "operations" means the number of 64-bit floating point operations needed to complete the indicated calculation, "memory" means the primary computer memory required to complete the calculation without resorting to significant use of secondary storage, and "data access" is the archival storage needed to save the results of the simulations (without compression).

The climate modeling effort at Los Alamos is a well-integrated initiative involving research and model development sponsored by the DOE HPCC Program and the DOE CHAMMP Program. Work sponsored by each program is mutually supportive and directed toward a common goal. The primary emphasis at Los Alamos in the CHAMMP Program is on ocean modeling, and the HPCC-sponsored activities are complementing that work and adding important new capabilities. Our specific future activities are focused on three issues involving mathematical and computational aspects of ocean modeling:

1. applications of differential sensitivity analysis to global ocean modeling,
2. a uniform finite-difference grid for a massively parallel global ocean model, and
3. barotropic closure models for large-scale ocean dynamics.

The ability to predict the mechanical response of new materials would dramatically enhance the economic competitiveness of the United States. Los Alamos is a unique national resource in both materials research and in computational capabilities. The Novel Materials Modeling Project will combine these two strengths to address a scientific problem that spans a truly wide range of physics—from the tiny time and distance scales of atoms and molecules to the next higher level of defect structures to the continuum scale of ordinary experience. Each level of theory has challenges of its own, namely, to incorporate the reality of the physics of scales below it as well as to provide physical guidance to the next higher scale. This latter challenge requires extraordinary computational resources. Our ultimate goal is to link these levels into a set of computational tools and to apply these tools to study the mechanical response of materials. Particular emphasis in FY92 has been placed on molecular dynamics simulations of ductile metals on the CM-200.

We are applying recently developed and tested computational algorithms (also referred to as Lattice Boltzmann methods) to fundamental problems in the simulation of multiphase flow in porous media in order to develop the basic theoretical and experimental foundation needed to advance this field. The Multiphase Flow in Porous Media Project has concentrated on developing better simulation capabilities for

Project Description	Operations	Memory (GBytes)	Data Access (TBytes)
Global Ocean			
century, 40 levels, 1/4° zoning	10^{17}	4	20
Porous Media			
3-D immiscible flow	10^{18}	1000	4
Novel Materials			
3-D molecular dynamics	10^{18}	20	3
3-D multimaterial hydrodynamics	10^{18}	1000	20
Plasma Physics			
numerical tokamak	10^{18}	1000	100

the processes involved in improved oil recovery (IOR), but the conclusions and methods to be developed will be equally applicable to the movement of immiscible liquids found in many hazardous waste sites. This work begins with research at the pore scale, and through existing collaborations with US petroleum industry scientists and engineers (e.g., AMOCO, MOBIL), proceeds to test the resulting methods and conclusions in laboratory and field-scale applications. The work is centered around the use of massively parallel computers (in particular, the Thinking Machines Corporation CM-2 and CM-5), and it integrates efforts already begun at other participating institutions (Lawrence Livermore National Laboratory and Notre Dame University).

SOFTWARE COMPONENTS AND TOOLS

One of the most crucial challenges of the HPCC Program is the visualization of results, intermediate and final, of data produced by the Grand Challenges. The size of the computations produces data sets of several gigabytes per recorded time step, which will overwhelm current visualization technology both in terms of hardware and software. Grand Challenge scientific models typically produce time-dependent spatially varying data. Visualization and data analysis require very-high-speed I/O to the high-performance computing environment (supercomputer, storage subsystem, graphics workstation). In addition to the existing hardware limitations, visualization must provide a means for data exploration of these massive data sets. Visualization tools will need to facilitate data perusal as well as provide new and useful methods of representation and presentation of data in a visual form.

There are several specific avenues of research that will fulfill these critical needs: high-level tools for scientific visualization and data exploration, new visualization algorithms, and distributed visualization. At the HPCRC, we are focusing on researching and developing such a visualization environment with each of the components incorporated in the GRACCE environment: CM-5, HPDS, and HIPPI MCN (discussed elsewhere in this report). Working with the Khoros Project at the University of New Mexico, we can leverage off their expertise in visual programming and distributed computing environments.

COMPUTATIONAL TECHNIQUES

An essential part of Los Alamos' Grand Challenge Computing Environment consists of supporting software for the high-performance systems in the Advanced Computing Laboratory (ACL). These software packages allow researchers to quickly solve difficult numerical problems with state-of-the-art algorithms, provide increased performance, and improve programmer productivity. We have contributed high-performance communications software to the Connection Machine Scientific Software Library (CMSSL), and we are currently working on further software and numerical algorithms for this library, particularly in the area of solution of sparse linear systems of algebraic equations.

Oak Ridge National Laboratory

GROUNDWATER CONTAMINANT TRANSPORT

The primary goal in this Grand Challenge is to apply the continuous increases in computational power and improvements in computational models to make significant advances in solving the problems of modeling transport phenomena and the ultimate fate of contaminants in groundwater. In FY92, efforts have focused upon "porting" an existing model to run on MIMD parallel computers for benchmarking and upon the development of flow and transport codes that will provide the basis for our new supercomputer model. FEMWATER, a groundwater code with considerable recognition in the community, has been converted to run on the Intel iPSC/860. This new parallel code, PFEM, has been tested on a simple problem and the results of testing and sensitivity analysis were presented. Validation of the model against data for a specific site on the Oak Ridge Reservation is in final stages. The universities funded under this program completed workplans for development of the new codes and code testing. Codes for saturated and unsaturated flow and transport with three-dimensional capabilities are currently under development. Additional studies involving identification of primary parametric uncertainties and data length scales for the selected site are underway. Periodic meetings are held to discuss workplans, compatibility of flow and transport codes, and current parallel programming research.

MATERIALS PROPERTIES

The central goal in the materials properties Grand Challenge is to develop and validate models for performing first principles density functional theory simulations of materials properties on complex systems. The task of the ORNL effort is to develop scalable parallel techniques and algorithms based upon multiple scattering theory capable of simulating systems with several hundred atoms. In FY92 we made progress in parallelizing existing methods and in generalizing them such that they can be used as test beds for new methods and codes and in testing some of the principles on which the code to perform large scale simulations will be based.

Research has concentrated on extending the parallelization of the Korringa-Kohn-Rostoker coherent potential approximation (KKR-CPA) method to include both tightly coupled and loosely coupled multiprocessors via PVM, generalization of the KKR-CPA method to include non-spherical contributions to the crystal potential, and the development of new parallel methods for simulating properties of large systems. Ideas in the latter area that have been investigated include using a large cell version of the KKR method and a new "local real space self-consistency technique." The latter technique has proved very accurate during tests on small systems.

NUMERICAL TOKAMAK

The computational science objectives of the ORNL Numerical Tokamak Project are: (1) the evaluation of CM-2, iPSC/860, Touchstone Delta, and Cray III computers and workstations using a fluid model and (2) the setting up of a common file format. On the first, the evaluation of the iPSC/860 at ORNL has been completed and concluded that the parallel algorithm for plasma turbulence calculations ran 5 times faster using all 128 nodes on the iPSC than a single-processor Cray-2. The evaluation of the Delta at Caltech gave results 16 times faster using all 512 nodes of the Delta. A banded convolution routine has been written and evaluated for these calculations for the CM-2. With the collaboration of JPL researchers, it has been possible to complete code optimization studies, that is, to find the optimal number of processors for a problem of a given size. From these results, we predict a factor of 100 speedup for the Intel Paragon machine. This can lead to a significant improvement in plasma turbulence calculations. In relation to the second objective, NASA's BDX machine-independent binary format has been found easier to implement in our code than netCDF, but netCDF stores information about the dimensions and locations of variables, which would make sharing data easier.

QUANTUM STRUCTURE OF MATTER

The Quantum Structure of Matter collaboration proposes a research program of forefront investigations on the fundamental structure of matter on the microscopic scale. Fundamental restructuring of the Hartree Fock codes for the Paragon has begun. Benchmarks have been established on the Intel iPSC/860 Hypercube for these codes, which are now running at speeds in excess of one Gigaflop. The collaboration has also begun operating a new Silicon Graphics workstation for high performance graphics enabling the visualization of computer-generated data.

SOFTWARE COMPONENTS AND TOOLS

Researchers trying to solve Computational Grand Challenges on high performance systems have common needs in software technology and programming environments, including software tools, parallelization tools, data management, and visualization. The overall goal of our software tools research is to improve the productivity of scientists working on Grand Challenges by first developing tools to assist them in writing programs for massively parallel computers and second to develop visual displays to assist them in understanding the massive amount of data generated by the programs.

In the tools arena, our Grand Challenges problems have relied on PICL for its portability and, along with ParaGraph, for performance monitoring and on PVM for distributed computing. In FY92 a new trace format for PICL was proposed to help handle the large amounts of data the Grand Challenge problems will create and for real time tracing, and the ParaGraph package has been updated to allow views of up to 512 processors. The PVM package has been supported and designs and specifications for the next major update due in FY93 were done. Research experiments in load balancing and shared memory using PVM were performed. A graphical interface tool called HeNCE, to make programming simple and intuitive for scientists and engineers, continued to be developed with a prototype version released for beta testing.

SUNY has concentrated on scientific volume visualization in general and output visualization of the materials properties research in particular. Atom positions and motions of colliding fullerenes from computer output generated at Ames have been visualized to illustrate the potential of their research. Texas A&M has continued to work on visualizing 3-D output from transport in porous media. Work on data compression using wavelets has been explored during FY92 in anticipation of the large data volumes generated by the groundwater research. Rice has continued to investigate the development of 2-D and 3-D interactive visualization tools. Early work on a visualization tool to monitor and control the progress of the simulation has begun.

COMPUTATIONAL TECHNIQUES

Computational techniques research supports the development of new parallel algorithms, numerical and mathematical analysis, and conceptual models of scientific applications. ORNL and its partner universities are performing these tasks with a focus on groundwater transport and material properties.

An improved version of the serial FEMWATER code for groundwater flow has been developed that is both faster and requires less memory than the original version. By the end of FY92 a PVM version of the FEMWATER code should be complete, which will increase the code's usefulness to researchers around the country. Research into preconditioners and iterative solvers for the groundwater model equations is just beginning.

Algorithmic improvements and fine-tuning changes in the parallel codes employing the KKR-CPA method for materials have been incorporated. ORNL researchers have investigated new methods applicable for distributed matrix diagonalization of a general dense matrix and have incorporated the algorithm into a materials science application code. Testing and enhancements will continue into FY93 when the Paragon arrives.

GLOBAL CLIMATE MODELING ON MP AND MIMD COMPUTERS

The goal of this work is to help develop the computational methods and tools needed to perform global coupled atmospheric/oceanic simulations. Finite-difference ocean models will serve as a platform for the development of improved solution algorithms, domain decomposition methods, and efficient linear equation solvers for MP computers, with applications to both atmospheric and oceanic models. We have completed an initial implementation of the Semtner-Chervin global ocean model for a two-degree grid test case on the nCUBE 2 which predicts time-dependent, 3-D velocity, heteropterous, and salinity fields in the multiconnected world ocean. Strategies for performing coupled ocean/atmosphere computations are being investigated. This work is being integrated with the efforts of other Laboratories in the development of DOE's advanced climate model. Sandia investigators will work closely with LANL CHAMMP researchers on the development of efficient parallel methods for simulation of global ocean circulation.

HIGH-PERFORMANCE ALGORITHM DEVELOPMENT AND ANALYSIS (HPADA)

The purpose of this task is the research, development, and analysis of innovative, high-performance algorithms for the numerical solution of partial differential equation-based Grand Challenge-class problems on advanced architecture computers. The initial applications focus is on issues which are common to both earth systems modeling and computational aerodynamics. In particular, the applicability and effectiveness of computational fluid dynamics solution methodologies transferred from aerodynamics to climate modeling will be investigated. A widely used ocean model will serve as a platform for investigation of issues related to mapping problems onto next-generation computers, including the use of highly efficient parallel linear system solvers; adaptive or embedded meshes; multigrid; and domain decomposition and load balancing techniques for irregularly structured meshes. A further important aspect of HPADA is the top-down analysis of the performance of the algorithms developed. To this end, the algorithms may be implemented on a variety of advanced architectures. This aspect of the work will be closely coordinated with the Computer Systems Performance Analysis Project managed by the Ames Laboratory.

ADVANCED COMPUTATION, COMMUNICATIONS RESEARCH, AND ASSOCIATED ACTIVITIES

HIGH-SPEED NETWORK RESEARCH AND MULTIMEDIA APPLICATIONS

This work is directed toward the identification and elimination of performance limitations in high-speed networks. The program uses network models and analytical tools to determine the practical limits of existing networks and develops solutions to these bottlenecks. Specific areas of interest include gateway and router congestion, TCP/IP enhancements, and resource allocation. The group has also developed a new approach to workstation interfaces.

A second component of this program is the development of tools for video and audio conferencing using conventional workstations. To provide such facilities requires development of tools to put audio and video into conventional packets. Since audio must arrive as a regular stream to be easily understood, we have developed algorithms to synchronize packets at the receiving computer. We have also developed a design for a whiteboard that can be shared over the network.

NATIONAL ENERGY RESEARCH SUPERCOMPUTER CENTER (NERSC)

The supercomputing environment at NERSC includes an array of services for users:

- General consulting
- Network information, support, and troubleshooting
- User training and education
- Documentation services
- Support of online bulletin boards
- A monthly newsletter (*The Buffer*) that focuses on systems issues and applications
- A library of applications code abstracts containing more than 300 entries.
- Consulting support for more than 115 applications codes.

We have available a number of important user libraries, and we actively support the multitasking of applications codes. We also provide advanced graphics software, such as the visualization tools produced at the National Center for Supercomputing Applications.

In FY91 the NERSC computer resources were distributed as follows:

Magnetic Fusion	39.69%
High Energy and Nuclear Physics	26.92%
Basic Energy Sciences	19.19%
Health and Environmental	10.73%
Superconducting Supercollider	3.47%

FY92 was characterized by profound change and redirection at NERSC. Conversion to a vendor-supplied operating system for the Cray computers was accomplished, mainframe hardware and software acquisition contracts were modified, and new strategic directions in parallel computing, distributed computing, and storage systems were defined. Discussion of progress and plans in these technical thrust areas is contained below.

TECHNICAL PROGRESS IN FY92

1. The largest effort completed this year was the conversion to the Cray Research Incorporated UNICOS operating system for the X-MP and two CRAY-2 supercomputers located at NERSC. As UNICOS was installed, we needed to allocate file systems, configure NQS batch queues, and tune scheduling algorithms. In addition, we implemented a dual-direction, automatic file migration between UNICOS and the Common File System to help alleviate the pressures on the online Cray disks.
2. An interim resource allocation and accounting mechanism was installed on UNICOS to provide necessary basic functionality. Work was begun on a centralized allocation management facility as requested by our sponsor and users.
3. We began acting as the operational site and supplier of scientific applications for the newly established National Storage Laboratory (NSL). We began collaborations with the IBM Federal Sector Division (system integrator for the NSL) to develop system management tools for NERSC's current and future production archival storage environment.
4. An interim CRAY Y-MP was installed at NERSC before delivery of the Y-MP C-90. In addition to the normal tuning/allocation work, we provided user access to its Solid-state Storage Device (SSD). This experience will also be useful when the C-90 arrives.
5. Continued exponential growth in the data generated by users (now in excess of 6 terabytes) required us to look for hardware relief for the manual shelf operation as well as a mechanism to transfer to a UNIX-based, high-performance, high-capacity storage system. Evaluation of successor storage systems continued.
6. We began analysis of various massively parallel architectures as potential follow-ons to the current generation of supercomputers.

INTERAGENCY INTERIM NREN

The Energy Sciences Network (ESNet) has established a lead role in encouraging and evaluating "cornerstone" technologies expected to be essential to the success of the National Research and Education Network (NREN) effort. These new technologies will enable data communications networking to evolve from the current T1-based bandwidth of 1.5-Mbit/s capacity to substantially faster speeds, with steps to 45, 155, and 622 Mbits in the next few years, giving a total speedup of more than 400 times the current bandwidth.

We have completed a competitive RFP, resulting in the selection of a contractor to provide broadband communication services to upgrade the ESNet.

NASA has worked closely with DOE in establishing this new technical direction and will use the same contract to similarly upgrade the NASA Science Internet (NSI). The National Science Foundation also expects to connect one site to gain experience with the new technology.

HIGH-PERFORMANCE COMPUTING RESEARCH CENTER

The Advanced Computing Laboratory (ACL) provides the foundation upon which the DOE HPCRC is built. The great majority of the hardware, software, and systems support for the HPCRC will be the responsibility of the ACL staff. The staff is composed of five basic groups: system scientists, administrative and technical support, computational scientists, computer scientists, and applied mathematicians.

The Thinking Machines Corporation Connection Machine 5 (CM-5) is the computational centerpiece of the DOE HPCRC. The CM-5 was delivered in late February 1992. We expect the system to be fully configured and fully functional by March 1, 1993. The basic parameters of the system will be 1024 nodes, 4096 vector processors, 32 Gbytes of memory, 200 Gbytes scalable disk array, four internal HIPPI channels, and eight Sun control processors. This machine, although compatible with slicewise software written for the CM-2, has several fundamentally different features. Specifically, the CM-5 supports both Single Program Multiple Data (SPMD) and Multiple Instruction Multiple Data (MIMD) styles of parallel computation in addition to low-overhead SIMD computations as on the CM-2; the CM-5 processor interconnection topology is completely different (fat tree) and designed to provide efficient communications for a larger set of applications; and the CM-5 relies on the Los Alamos-developed HIPPI for frame buffer visualization and debugging functions. We are continuing the close, synergistic relationship with Thinking Machines Corporation that began three years ago with the ACL's acquisition of the CM-2. For the next several years, the CM-5 will be one cornerstone of our Grand Challenge Computational Environment. Our goal is to move this enormous computational resource into the mainstream of production scientific computing.

GIGABITS RESEARCH AND DEVELOPMENT

The basic fabric of the Grand Challenge Computational Environment is provided by the Multiple Crossbar Network (MCN) Testbed. This is a high-performance local area network testbed at Los Alamos that supports bandwidths in the gigabit-per-second regime and utilizes emerging industry standards. Recent developments in high-performance networking technology offer the promise of interconnecting computational, data storage, and visualization systems at bandwidths two orders of magnitude greater than what most scientists have available today. This enormous increase in bandwidth will enable computational and data resources to be combined in new and different ways to facilitate the solution of Grand Challenge problems. Many Grand Challenge problems, such as global climate modeling and tokamak modeling, will generate terabytes of information. Gigabit-per-second networks and high-performance data storage systems are required to analyze the results of these calculations. Other Grand Challenge problems may benefit from the combination of supercomputer resources of different architectures linked together in a gigabit per second network. A gigabit-per-second network will be one of the enabling technologies for the Grand Challenge Computational Environment.

We are also working on a HIPPI-to-SONET (Synchronous Optical NETwork) gateway that will enable HIPPI networks to be extended using commercial optical fiber. This work provides technical support for national access to various high-performance computing resources such as the CM-5 or HPDS at Los Alamos.

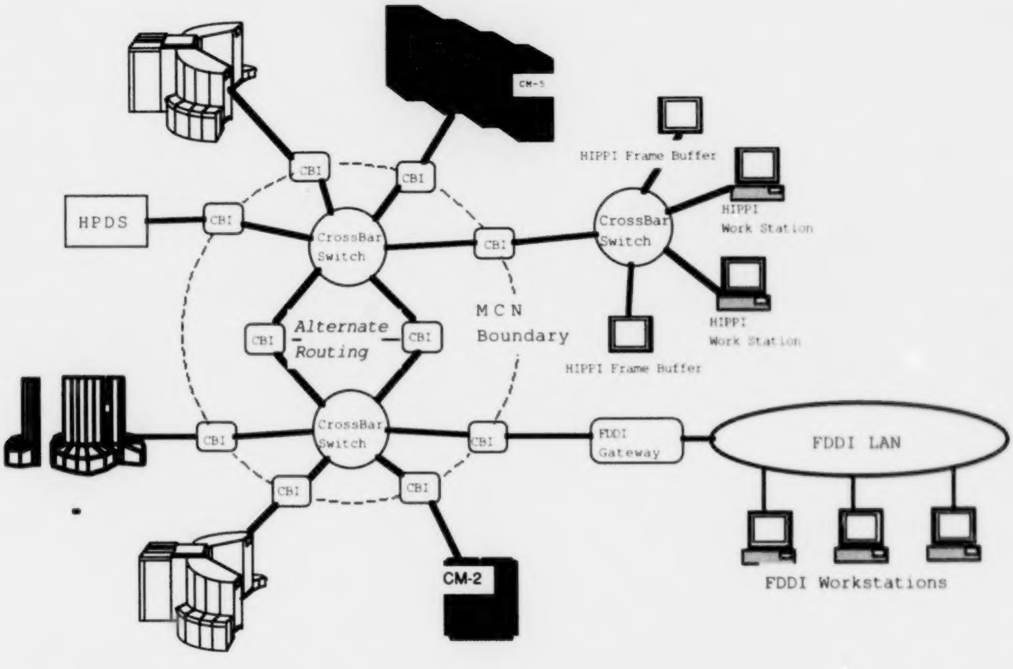

The MCN testbed links different architectures in a gigabit-per-second network.

HIGH-PERFORMANCE COMPUTING RESEARCH CENTER

The mission of the Computing Research Center component of ORNL's High-Performance Computing Research Center (HPCRC) is to provide the computing resources required to solve complex computational problems in various scientific disciplines that will benefit DOE and the nation, emphasizing groundwater transport, materials properties, and global climate modeling, and to support educational projects in the spirit of the HPCCP. Management support for ORNL's HPCRC will be provided by the Director of ORNL's Center for Computational Sciences (CCS), the CCS Industry Relations Manager, and the CCS Computer Resources Manager, as well as minimal professional and support staff.

FY92 accomplishments of the HPCRC include establishing a Cooperative Research and Development Agreement (CRADA) with Intel to evaluate beta versions of future Intel computers and acquiring the first beta system under this CRADA, finalizing an agreement with Intel to provide the HPCRC's massively parallel supercomputer, negotiating satellite systems to be placed at the selected sites of ORNL's Grand Challenge consortium, Partnership in Computational Science (PICS), and initiating the procurement of a second ring for the Kendall Square KSR-1 system.

NATIONAL RESEARCH AND EDUCATION NETWORK

ORNL will develop and maintain local networks appropriate to a High-Performance Computing Research Center. The objectives for this effort are ensuring that the networking capabilities of the ORNL HPCRC are adequate for researchers working on the various Grand Challenges, maintaining connectivity to external networks, and providing low-level implementation support for networking research in the Gigabit network project. FY92 NREN accomplishments center around the planning for a gigabit network between the two DOE HPCRCs (Los Alamos and Oak Ridge).

Sandia National Laboratories

COMPUTATIONAL DESIGN OF CATALYSTS AND BIOCATALYSTS (CDC)

The goal of CDC is to develop breakthrough capabilities in materials simulation using innovative new MP algorithms and methods and MP computers. The fundamental significance and broad economic, societal and scientific impacts of this project will be realized through a strong interdisciplinary approach involving partners with expertise in catalysis, chemistry, materials science as well as high-performance computing, and by coupling with experimental programs to validate the progress of computational results.

The partners are focusing on chemical and biocatalysis-catalysis for energy production, pollution minimization, and biotechnology. Some of our complementary efforts include micro- and nano-electronic and photonic materials for computers and consumer electronics, and novel structural materials for manufacturing. The computational techniques for these materials are closely related, thus allowing a coordinated attack on MP methods at three levels: first, development of scalable, microscopic quantum mechanical models (including Hartree-Fock, Local Density Approximation, and Many-Body methods); second, atomistic molecular dynamics and Monte Carlo methods for classical, microscopic force-field simulations; and, third, macroscopic continuum models.

Our vision is to be able routinely to do fully detailed 3-D simulations for the continuum models, microscopic force-field calculations with tens of millions of atoms, and full quantum mechanical relaxations for systems with hundreds and thousands of atoms. A nested capability is planned in which the continuum models provide boundary conditions for the force-field calculations and the force-field calculations are used to provide the boundary conditions for the quantum mechanical simulations. The full achievement of this vision requires development and availability of teraflop, terabyte computers.

GLOBAL CLIMATE MODELING ON MP, MIMD COMPUTERS

The goal of this work is to help develop the computational methods and tools needed to perform global coupled atmospheric/oceanic simulations. Finite-difference ocean models will serve as a platform for the development of improved solution algorithms, domain decomposition methods, and efficient linear equation solvers for MP computers, with applications to both atmospheric and oceanic models. We have completed an initial implementation of the Semtner-Chervin global ocean model for a two-degree grid test case on the nCUBE 2, which predicts time-dependent, 3-D velocity, temperature, and salinity fields in the multiconnected world ocean. This work is being integrated with the efforts of other Laboratories in the development of DOE's advanced climate model. Sandia investigators will work closely with LANL CHAMMP researchers on the development of efficient parallel methods for simulation of global ocean circulation.

HIGH-PERFORMANCE ALGORITHM DEVELOPMENT AND ANALYSIS (HPADA)

The purpose of this task is the research, development, and analysis of innovative, high-performance algorithms for the numerical solution of partial differential equation-based Grand Challenge-class problems on advanced architecture computers. The initial applications focus is on issues that are common to both earth systems modeling and computational aerodynamics. In particular, the applicability and effectiveness of computational fluid dynamics solution methodologies transferred from aerodynamics to climate modeling will be investigated. A widely used ocean model will serve as a platform for investigation of issues related to mapping problems onto next-generation computers, including the use of highly efficient parallel linear system solvers; adaptive or embedded meshes; multigrid; and domain decomposition and load balancing techniques for irregularly structured meshes. A further important aspect of HPADA is the top-down analysis of the performance of the algorithms developed. To this end, the algorithms may be implemented on a variety of advanced architectures. This aspect of the work will be closely coordinated with the Computer Systems Performance Analysis Project managed by Ames Laboratory.

Advanced Computational Methods for Rational Drug Design

The goal of this effort is to speed up the process of developing new drugs by complementing current discovery techniques, based on art and trial and error, with design techniques based on scientific principles. To that end, Sandia, in collaboration with Johns Hopkins School of Medicine, a major US pharmaceuticals company, and biomedical groups at other national laboratories and universities, is developing advanced computer modeling techniques that will couple experimental medicinal chemists with the tools of theoretical chemistry. This synergy between theory and experiment has already produced several important results, including a new understanding of the properties of the anti-cancer drug cyclophosphamide under conditions like those in the human body. Because of the large size of typical drug molecules and the complex chemical environment in which they work, new enabling technologies must be developed before all of our goals can be realized. (For example, the study of cyclophosphamide, a relatively small drug, involved the largest all-electron calculations ever performed.) The centerpiece of these new technologies is a large-scale parallel quantum chemistry package that runs on both massively parallel and distributed computers. Other new tools developed by Sandia include continuum-solvent methods for modeling aqueous solvation, genetic algorithms for optimizing chemical structures, and new acceleration methods for molecular dynamics simulations.

Ames Laboratory

BASIC RESEARCH AND HUMAN RESOURCES

MATHEMATICS OF PHYSICAL SYSTEMS

The applied mathematics effort at the Ames Laboratory is focused on wave propagation problems with special emphasis on direct and inverse scattering problems.

Novel methods of solving transient inverse scattering problems have been developed. These methods meld invariant imbedding and wave decomposition, or splitting, methods to yield physical insight into the underlying scattering processes. At the same time the methodology suggests effective algorithms for the numerical solution of both the direct and inverse transient scattering problem. Additional work is carried out in the asymptotic analysis and solution of scattering problems. Acoustic, elastic, viscoelastic and electromagnetic problems with applications to nondestructive evaluation, geoprospecting, medical imaging and subsurface contamination identification have been studied. Program members are involved in extensive collaborative efforts with university faculty in the United States, Great Britain, Sweden and New Zealand. The program also enjoys an ongoing collaboration with the USAF School of Aerospace Medicine.

COMPUTATIONAL SCIENCE FELLOWSHIP PROGRAM

The purpose of this program is to increase the number of professionals in the area of computational science. The program provides graduate fellowships for up to 60 graduate students in any one year. Fellowships are extendable for up to four years. The program elements include

1. active solicitation of applications for the program
2. selection of fellows with the assistance of an advisory panel
3. selection of participating universities with the assistance of an advisory panel, and

4. matching of students and DOE sites for a three-month practicum during the tenure of the fellowship. The program provides tuition and fees, support to the host department for the direct benefit of the student, and the cost-sharing of a work station to be used by the student.

CURRICULUM ENHANCEMENT PROJECT

The purpose of this project is to encourage and support activities that promote change in university curricula that result in broader acceptance of Computational Science as a discipline at the graduate level. Activities include

1. developing and supporting experimental course materials,
2. providing a forum for faculty and administration to share experience and strategies for implementing computational science curricula, and
3. encouraging and supporting articles and other broadly distributed materials about the value of computational science.

Argonne National Laboratory

APPLIED MATHEMATICS AND COMPUTER SCIENCE

Mathematics of Physical Systems

Our research centers on methods and algorithms for the study of vortex lattices in high-Tc superconducting materials. Investigations are based on macroscopic (phenomenological) models—in particular, the London and Ginzburg-Landau models for weakly anisotropic materials and the Lawrence-Doniach model for highly anisotropic materials. A major accomplishment has been the development and implementation of special optimization and linear algebra techniques for the minimization of the Lawrence-Doniach free-energy functional. An inexact Newton method enables the computation of vortex configurations with hundreds of vortices per unit cell in layered superconductors with several layers. A reordering of the matrix based on a parallel graph coloring heuristic is used to achieve scalable parallelism.

Optimization and Mathematical Programming

Our work on large-scale optimization algorithms centers on the development of a trust region method for large-scale optimization problems. One of the main ingredients in the trust region method is the incomplete Cholesky factorization, developed by Argonne's linear algebra group for the solution of sparse positive definite problems; we are using this factorization for the solution of the indefinite trust region subproblems encountered in optimization problems. The combination of the trust region method with incomplete Cholesky factorizations and with graph coloring algorithms for the estimation of sparse Hessian matrices has resulted in a truly powerful optimization code. We are also investigating the optimal control of systems described by ordinary differential equations, using an interior point methodology to solve problems with mixed state and control constraints. We believe that this methodology will prove to be a rich source of practical algorithms for solving a new class of optimization problems known as extended linear-quadratic programming problems.

Geometric and Symbolic Computation: Automated Reasoning

Logical reasoning is required for many important problems in science, for example, electronic circuit design, real-time systems control, and program verification. Automated reasoning programs offer scientists powerful capabilities for attacking such problems: the ability to search for theorems, construct mathematical objects, test hypotheses, verify the correctness and completeness of needed proofs, and supply proofs in situations where scientists had failed for years in the corresponding attempt. We have developed the world's most powerful general-purpose reasoning program, OTTER. Important to OTTER's usefulness in attacking complex scientific problems was the development of a new indexing technique, which enables OTTER to have almost linear performance over time. We have also formulated the linked inference principle, which sharply reduces the retention of unnecessary clauses. These advancements have enabled us to answer open questions in logic and mathematics.

Numerical Analysis and Scientific Computing: Linear Algebra and PDEs

A major focus of the linear algebra group at Argonne has been the development of communication-minimizing algorithms. We have taken two approaches: domain decomposition methods and incomplete matrix factorization (IMF) methods. Domain decomposition methods seek to minimize global communication by iterating between two problems: a set of independent local problems and a small global problem. The local problems require no communication, and little communication is required for the global problem because it is small. In our IMF methods we minimize communication by using a graph-based reordering of the matrix to reduce the number of messages sent by each processor. Our focus for both approaches has been on sparse linear algebra problems arising from physical phenomena. For example, we have looked extensively at static displacement and vibration analysis problems arising in structural mechanics. A second aspect of our research is the design of algorithms for sparse problems with a specific structure.

In particular, we are developing algorithms that can exploit the predictable structure in problems arising in the areas of ordinary differential equations and nonlinear optimization.

Systems Software

Many scientific computations rely heavily on basic numerical software, such as the elementary and special functions. We are developing highly accurate elementary function software for all of the commonly used functions. We are also developing software technology that will permit the automated implementation of these new functions and algorithms on diverse high-performance architectures. To ensure the accuracy of the implementations, we are designing rigorous error analysis techniques. These efforts are being conducted in collaboration with computer vendors (in particular, SUN), through whom we plan to test and distribute our parallel software.

Data Handling and Analysis

The Database Computing Project is prototyping a scalable, extensible, object-oriented database designed to analyze high energy physics data via simple queries. The results of these queries can then be visualized by using a variety of tools that make it possible to locate the few events indicating the presence of a new particle from among hundreds of thousands of less interesting events. The user can interactively make "cuts" to produce smaller samples of events, and then interactively search through these samples with visualization tools to locate the particles of interest. Further, the user can walk through fairly complicated events consisting of hundreds of attributes.

Advanced Visualization Concepts

Large-scale parallel computing applications require powerful data analysis capabilities, including the ability to interactively explore massive datasets and display data distributed on parallel architectures during run time. Currently, the primary method for incorporating scientific visualization into applications has been through the use of tools operating on data in batch mode. We are developing a system of concurrent, interactive, real-time scientific visualization tools to monitor the results of, and interact with, massively parallel computations during run time, and to interactively navigate large terabyte-sized datasets. Preliminary experiments have included demonstration of the application of such tools to large-scale computational problems in computational biology, global climate modeling, and materials science.

Brookhaven National Laboratory

ANALYTICAL AND NUMERICAL METHODS

This is the longest established mathematical sciences program at Brookhaven. The major emphasis is on the development of methods for the efficient solution of the discrete systems resulting from numerical approximation of partial differential equations. This is pursued in two principal modes. First, there is work to advance the fundamental understanding of preconditioners and iteration. The importance of this lies in the fact that the increase in computational power over the past decades has led to an enormous scaleup of the size of problem which can be numerically tractable. There may be numerical methods whose errors or instabilities are insignificant for relatively modest problems but will become critical for large ones. For this reason critical mathematical analysis is necessary to develop confidence in methods which are used. Second, the work involves the search for faster and more robust algorithms for parallel computing architectures. The tools that are employed are functional analysis, regularity theory, and finite element error analysis. The emphasis is on exploiting the analytic structure of the partial differential equations underlying the numerical approximations, rather than concentrating solely on the algebraic properties of those approximations. Domain Decomposition and multilevel methods have been of particular concern. A second area of emphasis has been the development of algorithms for particular classes of problem, such as unbounded domains and singularly perturbed problems.

NONLINEAR DYNAMICS

The research in nonlinear dynamics focuses on the qualitative behavior of prototypical low-dimensional dissipative systems, such as Duffing's equation, which models, for example, the forced mechanical vibrations of elastic structures, or electrical oscillations involving a saturable core inductor. Methods from the geometric theory of phase space are combined with interactive numerical simulation and computer visualization techniques. Work also includes the application of dynamical systems approaches, especially time series analysis methods, to the question of climate variability and detection of the greenhouse signal.

Using computer visualization to conduct exploratory surveys by varying parameters in the equations has led to the discovery of many new phenomena. Of particular concern in applications are global bifurcations causing catastrophic loss of stability (blue-sky catastrophes or boundary crises). Recently discovered phenomena include chaotic saddle catastrophes, catastrophes with indeterminate outcome, codimension-2 bifurcation patterns, and a technique for detecting incipient catastrophe from analysis of chaotic time series data. This last topic involves analysis of experimental or empirical data without reference to a derivable mathematical model.

A recent discovery of importance to time series analysis is the fact that a fundamental invariant of a chaotic attractor, its Euler characteristic, is constrained by the index theory of fixed points. This is an integer invariant which is intrinsically more robust than noninteger invariants such as dimension and Lyapunov exponents, and it is the only known integer invariant that generalizes readily to higher dimensions.

FLUID MECHANICS: NEW MODELS AND NEW ALGORITHMS

A major component concerns the development, analysis, implementation and application of vortex methods. Vortex methods, which have proven to be key techniques for computing and analyzing turbulent flow, were created and developed within our group under DOE support, and can be used to understand the fundamentals of turbulent flow. Their popularity is due to two fundamental advantages over traditional techniques. First, they avoid the use of meshes, and hence neatly sidestep all the grid-generation problems that can overwhelm such techniques. Second, since they are Lagrangian techniques that automatically focus computational resources where necessary to resolve sharp gradients in the flow, they offer the chance to focus on the real fundamental processes that control the development and evolution of turbulence. Since their creation in the late 1970s, they have found their way into a large number of applications, including flows past submerged bodies, in valves and pistons, in air-cooling systems, past wings, in mixing chambers, and as a component of combustion calculations.

Within the LBL Mathematics Department, vortex methods have been used to study turbulence and critical phenomena in superfluid helium, to analyze vortex motion in the plane, and to understand combustion and the stability of flames. This year the work in vortex methods has been extended to study fully three-dimensional turbulent flow, and expanded to include some of the small-scale turbulence issues that reveal the connection of fluid motion to statistical mechanics.

To begin, a fundamental aspect of three-dimensional flow is the development of thin vortex elements and filaments resulting from vortex stretching. The correct representation of small scales of this motion is imperative in order to maintain the correct energy and momentum of the flow. An appropriate technology to model the delicate pinching and cusping of hairpin fluid structures has been developed, and algorithms have been devised to replace excessive pinching with the correct aggregate effect. At the same time, a new, highly accurate technique relying on vortex loops has been developed for computing the evolution of vorticity off solid walls.

These techniques have come together in the development of a working code to model three-dimensional viscous, incompressible fluid flow past arbitrarily complex bodies. The code combines a vortex cylinder method, together with a panel method for satisfying the noslip condition. This code, fully vectorized and running, is currently being used to model wake development, eddy shedding, mixing, transport, far-wake evolution, and fluid-body interactions. Visualized using the real-time visualization environment developed by our group for massively parallel processors, our results document the intricate mechanisms of vortex shedding, stretching, and rollup behind complex bodies. As a test case, detailed studies of flow past a sphere at a variety of Reynolds numbers have been performed and produce accurate prediction of the drag coefficient, and some preliminary evidence that suggests the ability to predict the onset of the drag crisis. This code has also been used to study flow past airfoils, stall effects, and flow response to upstream vortex structures.

Other areas of work concerning vortex methods include analysis of wave propagation and solitary waves on interacting vortex filaments, study of long-term vortex breakdown, analysis of a system of magnetization variables associated with a vorticity perspective, modeling of ring and shear calculations, and linking of vorticity equations to conservation laws for vorticity and associated entropy conditions.

The technology of vortex methods has been employed in a variety of applications outside of the LBL work, and collaborations have been established, at various times, with efforts at the University of California, Berkeley; University of California, Los Angeles; Massachusetts Institute of Technology; Princeton University; Lawrence Livermore National Laboratory; Sandia National Laboratories; International Business Machines; Thinking Machines Corporation; the Naval Underwater Systems Center; Oak Ridge National Laboratory; and Schlumberger-Doll. Within LBL in the next two years these techniques will be carried in two directions. As theoretical tools, they will be instrumental in understanding the fundamental mechanisms behind singularity development in evolving flows, energy cascades, and the formation and persistence of large-scale global vortex structures that govern such diverse phenomena as meteorological flows, ocean currents, mixing and transport strategies, and vehicle dynamics. As practical computing tools, the marriage of vortex methods and massively parallel processors has resulted in a general algorithm for computing flows in a wide variety of engineering situations, and will be used to perform vortex simulations of flows past vehicles and full cycle simulations of an internal combustion engine.

SURFACES AND INTERFACES

Some of the most interesting problems in fluid mechanics, material sciences, and physics involve the position and propagation of interfaces between substances. Such phenomena arise in problems of crystal growth, flame propagation, Hele-Shaw cells, the Rayleigh-Taylor and Kelvin-Helmholtz instabilities, capillarity theory, flow in porous media, lithography, and device processes. Many of the problems are linked by common mathematical and numerical problems: the tendency of fronts to develop cusps and corners, pronounced topological changes such as breaking and merging, front propagation speed driven by local geometric properties such as curvature, and intricate feedback between the front configurations and the physics of the underlying media.

The LBL group has introduced a set of techniques for following moving interfaces that has opened up a spectrum of previously intractable problems. These techniques, known as level set algorithms, have the tremendous virtue that they work in any number of space dimensions and follow the most intricate interfaces, even as they develop sharp corners, singularities, and change topology. The central mathematical idea is to view the interface as the particular level set of a higher dimensional function, and then follow the evolution of the resulting initial value partial differential equation by exploiting technology borrowed from the numerical solution of hyperbolic conservation laws. These techniques have been used by such groups as combustion scientists to compute optimal burning configurations in flames, material scientists to compute dendritic growth in crystalline melts and fluid dynamicists to compute oil-water interfaces. At the same time, they have been altered to answer a series of theoretical questions about singularity developments in collapsing bubbles and thin films, and used to generate minimal surfaces through complex bounding frames.

This year the focus has been on applying this level set front tracking approach to a collection of physical problems. To begin, a general algorithm has been as devised to compute minimal surfaces passing through arbitrarily complex bounding frames, and used to examine the breakage in soap bubbles and thin films as they move. The level set technology has been coupled to projection methods for fluid flow to model the motion of fluid droplets and interfaces under surface tension; such problems are important in understanding cavitation and two-fluid interface dynamics. As a test problem, the difficult problem of three-dimensional fluid droplet falling into a bath has been studied.

In another application, the motion of complex solid/liquid boundaries in crystal growth has been modeled, including the physical effects of crystalline anisotropy, surface tension, molecular kinetics and undercooling. The results show the evolution of complex crystalline shapes, the development of large spikes and corners, and dendritic formation and side-branching. In collaboration with Oak Ridge, we have used the level set technique to study the evolution of bubbles in a prototype mixing device, and we have worked with computer scientists at the University of South Florida to design algorithms to measure tumors in medical imaging.

Outside of the LBL group, these techniques are being considered for use in a variety of problems, including image processing, the development of efficient inkjet plotters, flame calculations, crystal growth, control theory, combustion studies, secondary-oil recovery, and cavitation problems.

On a different front, the LBL group has spent considerable effort on a long and productive project to obtain fundamental understanding of capillary phenomena. This work is critical in application areas such as thin film formation, surface tension phenomena, material processing, and flow in porous media. It is a broad effort that goes from rigorous mathematical proof to numerical computation to physical experiment on space flights in microgravity and back. For example, the LBL group has made extensive analytical and numerical results that uncovered a class of "exotic" rotationally symmetric containers. One such container—which differs only locally near the mid-height from a circular cylinder when half filled with fluid in zero gravity—admits an entire continuum of symmetric free surface interfaces, including the horizontal one. All of these interfaces are stationary but not minimizing for the energy functional. There exists a minimizing surface, but it can be shown to be asymmetric, even though the container is symmetric. Such a container can be characterized for any given gravity field and contact angle, but it becomes very small, and effects would be difficult to observe, except under microgravity conditions.

The results are gaining increasing influence, and new opportunities are becoming available to the LBL group in the US and European space programs (NASA and EAS, respectively). The unique information provided by experiments in microgravity, not available in a terrestrial environment, permit physical corroboration of our mathematical and computational findings and provide realistic input for verifying the underlying physical concepts of contact angle and other aspects of the classical Young-Laplace-Gauss capillarity theory.

STATISTICAL MECHANICS, POLYMERS, SUPERFLUIDS, AND MATERIAL SCIENCES

The relationship between classical statistical mechanics, field theory, and turbulence has been insufficiently studied in the past because of the lack of a coherent approach that takes into account the properties of the equations of fluid mechanics. We believe that we have a new and coherent approach that allows us to use in fluid mechanics (including superfluid mechanics and eventually superconductivity) some of the new and powerful tools of statistical mechanics. At present, this effort is more theoretical than our usual efforts, but it has already led to practical computational tools.

In the past few years, we have constructed Gibbs ensembles of vortex filaments and showed that stable equilibria have a polymeric structure, maximum entropy, and a Kolmogorov spectrum; produced a clean derivation of the Kolmogorov law; derived an Ansatz that had been used *ad hoc* for the calculation of critical exponents in superfluidity; constructed the inertial range of turbulence by means of percolation theory and an explicit model of the interaction of superfluid vortex lines with their spin-wave background. In the classical turbulence problem, our results agree well with numerical results and help to explain them; in the quantum case, the results are in excellent agreement with the appropriate experimental data.

A major component of the analysis that has led to these results is the relation we discovered between the critical state of an ensemble of vortex filaments and the perimeter of percolation clusters in a vortex percolation problem. This relation makes it possible to characterize efficiently for the computer both a vortex system in fully developed turbulence and a quantum system near the lambda transition. By searching for all percolation states in the temperature/chemical potential parameter plane, we have been able to show that these states can be connected by a universality curve along which certain exponents are invariant. The existence of this connection has allowed us to translate the quantum mechanical real-space renormalization group from the quantum XY model to fluid mechanics; this renormalization is being used at LBL in the context of vortex methods.

There is another case where a transition for a system of filamentary objects can be characterized by a percolation property: a system of polymers in the plane with near-neighbor attraction, where the transition is between an unfolded and a folded state. Qualitatively, the vortex and polymer transitions are very similar. We have started a numerical investigation to see how general this type of transition is by looking for new examples. This investigation is likely to result in new understanding of fully developed turbulence as well as of polymer folding. In addition, the relation between vortices, polymers, and percolation is of substantial mathematical interest.

We believe that this effort provides a promising avenue toward modeling turbulence in classical fluids and understanding of the behavior of quantum fluids. In particular, the theory permits the analysis of the ways in which numerical calculations fail to converge in a turbulent regime. Most importantly, the analysis leads to a real-space renormalization of vortex interactions that can greatly simplify calculation in fluid mechanics at a very high Reynolds number (the turbulent case). Our methods also provide a new and unexpected methodology for studying quantum vortices and their interactions with the surrounding medium (in particular, with spin-waves, and, in the superconducting case, with a phonon background). The analysis has mathematical implications for the statistical mechanical theory of polymeric systems, and may well lead to results of practical usefulness in that context as well. In addition, our methods have similarities to recent constructions in field theory, and the possibilities of fruitful interaction between turbulence theory, the theory of quantum fluids, and field theory appear endless.

COMPUTATIONAL MATHEMATICS

Much of the work on material sciences concerns noninvasive reconstruction of materials through attention to eigen-functions and normal mode placement. The goal of this work is to detect the presence of cracks, defects, and aberrations in material elements such as beams and solid bodies by looking at how the material responds to oscillations. By considering both the positions of the nodes of the eigen-functions and their values, an efficient numerical algorithm has been devised that can determine, at least for some simple model problems, the density distribution of the material under study. These techniques can detect, recover, and reconstruct materials with discontinuous density and discontinuous modulus of continuity. Upper and lower bounds for the eigenvalues have been produced, and numerical experiments show that this bound is optimal.

We have also attacked eigenvalue problems for two-dimensional Schroedinger operators on rectangular domains. Around the corner is the good possibility of attacking inverse nodal problems for Schroedinger operators on two-dimensional domains. Here we expect to be able to show that the nodal lines for the eigenfunctions determine the potential uniquely. Just as in the previous case, we expect to use our new techniques to develop an efficient numerical algorithm for determining the potential from nodal data.

On the linear algebra side, much of our work has been on preconditioned iterative methods for large, sparse systems of linear equations. Recently our interests in iterative methods have been on systems that need not be positive definite. Such systems, real and complex, arise in many applications such as electromagnetics, fluid mechanics, and geophysical exploration. However, numerical methods for these problems are not nearly as developed as for the symmetric positive-definite case.

COMBUSTION AND ENGINEERING FLUID MECHANICS

The work here is on computing real engineering flows which arise in devices. Some of the fundamental workhorse codes of the engineering community—for example, projection methods for viscous flows, which lie at the core of the Livermore's CFD HPCC work—were created and perfected by members of the LBL Mathematics Department. These codes are perfect for low and moderate Reynolds number viscous flows and, together with vortex methods for high Reynolds number turbulent flows, provide a spectrum of computational techniques for accurately computing the full range of complex viscous flow phenomena. For example, we have coupled these methods together with combustion models, mixing transport techniques, interface methods, and shock schemes to study such problems as flows inside ramjet combustors, internal combustion engines, inkjet plotters, and wing designs.

Several engineering/computational fluid dynamics problems are of particular focus within the LBL group. To begin, combining vortex methods, elliptic solvers, and the level set Hamilton-Jacobi approach for tracking interfaces has yielded a full code for simulating flame propagation problems that arise in combustion systems. For example, this code has been used to study free-standing and attached flames; in the case of attached flames, we have made detailed experiments of the stability of a flame with respect to the various physical parameters, and shown the dependence of the flame brush on exothermicity rates, the cusping of a flame resulting from the influence of incoming turbulence levels, the widening of a flame brush as a function of turbulence levels, and the self-cusping of a flame as a function of curvature stretch-induced flame vorticity.

As another example, a code for modeling fully three-dimensional viscous compressible flow using vortex methods has been completed and is now in use to study the effect of body geometry and Reynolds number on drag profiles, eddy shedding, mixing statistics, and energy transport. A two-dimensional version of this code was used to performed a detailed study of the flow past a partially opened valve in a piston. This design geometry, suggested in collaboration with an automotive engine manufacturer, shows the development of large eddy structures behind the piston, the stagnant vortices in the main body of the piston chamber, and the turbulizing effect they have on the mixing process.

In the next few years, the LBL group hopes to develop hybrid algorithms that glue together tools that are under development in our group. The marriage of projection and vortex methods to level set interface techniques, fast algorithms, and massively parallel processors should allow us to tackle hard problems in solid/liquid melts, complex combustion, multifluid-fluid problems, stratified density flows, and meteorology.

PARALLEL PROCESSING AND ADVANCED VISUALIZATION

The group at LBL continues its work on the frontier of research into the development of realistic schemes for complex phenomena on massively parallel processors. Rather than focus on kernel applications, the focus of this effort has been the implementation of highly involved schemes onto state-of-the-art architectures, the evaluation of parallel processing as a truly competitive and possibly superior programming mode, and the creation of software advances required to bring parallel computing to routine programming environments.

The LBL group has been involved in a multiyear collaboration with Thinking Machines Corporation, and the goal of this work has been two-fold. First, to demonstrate the feasibility of massively parallel computing, both SIMD and MIMD, to practical, large-scale computing efforts. Second, to serve as a benchmark testbed environment during the architectural design stage and software development of the next generation of high-performance computers.

The work on Connection Machines has been used as a showcase demonstration of the power of such new architectures, including codes to model flow past piston valves, air-cooling flows in computers, and real-time interactive flow visualization.

The centerpiece of this work has been a two- and three-dimensional vortex method for computing viscous, incompressible flow, which requires efficient implementation of such issues as non-nearest neighbor topology, adaptive creation of computational elements, expanding and contracting data structures, and non-obvious load balancing. This code has been used to compute flow in a variety of regimes, and kernels of this code were used as a prototype test for the CM-5 as it was developed.

Finally, the LBL group has continued its work on real-time visualization, producing a flow environment on the Connection Machine that allows interactive injection of passive tracers, bubble wires, and ribbons that mimic experiments performed in experimental apparatus.

Lawrence Livermore National Laboratory

NUMERICAL METHODS FOR LARGE-SCALE SYSTEMS

This research effort develops and numerically analyzes advanced methods for solving large-scale ordinary differential and/or algebraic systems of equations, nonlinear optimization problems, and specific partial differential equations. The approach combines mathematical analyses of stability, approximation, and convergence with numerical tests on a variety of computational platforms (especially massively parallel architectures) to ensure high performance and reliability in DOE applications requiring such algorithms.

The FY92 accomplishments are substantial. Extensions of previous analyses of conjugate gradient algorithms to orthogonal error algorithms and complex linear systems unify the theory underlying these methods and problems. This work clarifies the relationship between seemingly unrelated issues and enables the development of more powerful methods. Studies on the solution of the linear Boltzmann equation have produced a linear algebraic formulation of a common class of Boltzmann equation discretizations. This technology has been used to analyze the asymptotic behavior of the associated linear operators in certain physical regimes arising in applications, particularly those involving radiative transfer. These results show why the preconditioner used in the Diffusion Synthetic Acceleration method is effective under very general assumptions that are easily satisfied in practical applications. The development and analysis continue on multigrid methods for the solution of Poisson type PDE problems, for which multigrid methods are well known to provide very efficient solution algorithms. A detailed analysis of multigrid methods as subspace methods has begun. Development and testing of algorithms for underwater acoustics have also begun. Preliminary versions on the BBN Butterfly and Intel Touchstone computers exploit the parallelism that arises from the near independence of the solution on different ray paths found in the normal modes approach.

LARGE-SCALE SCIENTIFIC COMPUTATION

The LLNL research program in Large-Scale Scientific Computation focuses on the development of high-resolution finite difference methods for fluid dynamics and on the application of these methods to the discovery and analysis of fundamental fluid dynamical phenomena. The approach to designing methods, based on higher-order extensions of Godunov's finite difference method, builds the dominant physical behavior of the solution differential equations into the algorithm. These discretization techniques combined with adaptive methods focus computational effort to obtain optimal performance on target applications. These adaptive techniques include local adaptive mesh refinement (AMR) and front tracking.

Three FY92 accomplishments stand out. The first is the development of a 3-D AMR algorithm for gas dynamics and other systems of hyperbolic conservation laws (Joint with CIMS and LANL). Tests of the method indicate that it reduces computational costs by more than an order of magnitude compared to fixed grid computations of comparable resolution. Second is a new volume-of-fluid technique for tracking shock waves that provides a robust framework that automatically treats highly irregular, kinked fronts and changes in frontal topology (jointly with UCB). This methodology also models complex problem geometry by representing domain boundaries as stationary "fronts" embedded in a Cartesian mesh. The third result is the development of a 3-D projection algorithm for the incompressible Navier-Stokes and Euler equations. This method has been successfully applied to the study vorticity intensification and the formation of "hairpin" vortices during the transition to turbulence of a co-flowing jet.

Los Alamos National Laboratory

RESEARCH PARTICIPATION AND TRAINING

The ability of the US to compete in world markets increasingly depends upon our ability to educate scientists, and an important part of that education is technology related. The Computational Science Workshop Series addresses this challenge. It provides professionals and students the opportunity to become immersed in workshops that are designed to provide advanced education in high-performance computing as it applies to computationally intensive research.

The workshops specifically address the three primary components of high-performance computing: methodologies, environments, and applications. Emphasis is placed on enabling participants to effectively pursue specific areas of interest through applied research and collaboration. Collaboration between participants and Laboratory scientists is encouraged through the workshop seminars and mentorship programs as well as through the Laboratory's numerous technical colloquia. Participants are also encouraged to pursue their individual research interests.

EDUCATION, TRAINING, AND CURRICULUM

The New Mexico Supercomputing Challenge is an academic-year competition that was initiated in 1990. Now in its third year, the Challenge is a program that allows high school students to do computational science projects using high-performance computers. The Challenge is open to all New Mexico students on a nonselective basis. The purpose of the Challenge is to expose students and teachers to computational subjects and experiences that they might otherwise not have. The emphasis is on achievement, or competition, at the teams' own level. Through a partnership of business, universities, and national laboratories, teams of students from public and private New Mexico schools have the opportunity to learn about and actually run programs on high-performance computers.

MATHEMATICS OF PHYSICAL SYSTEMS

The emphasis of the program in Mathematics of Physical Systems is on combining novel mathematical theory with state-of-the-art computation.

Kinetic equations arise in many scientific contexts, and because of their high dimensionality, direct numerical solution of kinetic equations is frequently impossible, even with the most advanced supercomputers. By using techniques known as "R-singular perturbations," we are able to reformulate general kinetic problems as two simpler problems: the inner problem of the boundary layers and the outer problem of regions away from the boundary layers. The significance of the reformulation is that it simplifies kinetic problems enough to make numerical solution practical. In particular, we are solving the kinetic equations to obtain a new, highly accurate model for semiconductors.

A related area of research is particle transport. Recently, we have developed massively parallel techniques for solving two distinct models—one a discrete linear approximation and the other a simplified spherical-harmonics approximation—for transport equations. Surprisingly, we have been able to establish a similarity relation between the two approximations and thus have been able to solve the spherical-harmonics model using a modified version of our massively parallel linear solver. Both codes are now running on the CM-2.

In the study of the properties of novel electronic materials, an important tool is provided by the methodology of quantum simulations. We succeeded in developing the first quantum Monte Carlo code to simulate the degenerate, single impurity Anderson model and implemented this code in a distributed computing environment, achieving Y-MP performance on clusters of 5 to 10 workstations. The code is currently being transferred to the CM-5 computer.

For another class of statistical mechanics problems, a key mathematical property, called the Markov property, has been used to develop new, specifically massively parallel Monte Carlo computational methods. The scientific problem under study is the validity of the fundamental, renormalization-group hypothesis of modern critical phenomena theory. Recently, a massively parallel Monte-Carlo procedure was developed to investigate a classical problem, namely, the renormalized coupling constant for the one-dimensional Ising model.

In the field of fluid dynamics, we are studying acceleration-driven fluid instabilities. We have developed theoretical models of the mixing layers for Richtmyer-Meshkov and Rayleigh-Taylor instabilities and are validating them by comparison with experiments and computations.

DYNAMICAL SYSTEMS AND CHAOS

The aim of the program is to develop techniques for identifying and studying behavior that appears to be highly complex or random, and yet is generated by a simple underlying mechanism.

Cellular automata (CA) are a novel class of simple mathematical systems defined on lattices of nodes, or processors. Practically, CA serve as models of the type of complex behavior generated by the interaction of many simple, locally connected components—such as lattice gases and percolation through porous media—and in the design of parallel architectures and general parallel-processing algorithms. Recent research has led to the solution of certain ostensibly "chaotic" CA. The solution method maps these systems onto an exactly solvable linear automaton. The solution method suggests the possibility of increased efficiency in implementing CA through dynamic allocation of processors.

Work in nonlinear time series analysis is motivated by the fact that apparent randomness may be due to chaotic behavior of a nonlinear but deterministic system. In such cases, it is possible with nonlinear analyses to make short-term forecasts much more accurate than is possible with a linear stochastic model. We are developing methods both for short-term forecasting and for identifying low-dimensional chaotic behavior. We are applying the method to various data sets—including sun-spot cycles, river and sea levels, measles and chickenpox epidemics, and EEGs—for which it has been proposed that the underlying mechanism is that of nonlinear chaotic dynamics.

In the study of nonlinear optics, the emphasis is on self-induced transparency (SIT), a phenomenon that arises when an optical pulse propagates resonantly through a nonlinear medium without loss. During the past year we used averaging and multiple time-scale techniques to obtain a new derivation of the SIT equations from the fundamental Maxwell-Schroedinger equations. In addition, we discovered analytically that the equations can exhibit chaos in two physically relevant situations: a single-mode ring-cavity laser perturbed by a low-intensity probe laser and a multi-mode ring-cavity laser with special couplings of dynamics among the modes. This chaotic behavior manifests itself as "flickering" observed in the light intensity emitted by the perturbed ring-cavity laser.

NUMERICAL ANALYSIS AND SCIENTIFIC COMPUTING

The program objective is to develop novel computational techniques for problems including numerical solution of partial differential equations (PDEs) and signal reconstruction from data.

Adaptive Mesh Refinement (AMR) is a numerical technique that incorporates nested, logically rectangular meshes to resolve phenomena too expensive to compute with uniform meshes. Our algorithms define a performance index that quantifies how well the mesh represents the solution and the size of truncation errors introduced by discretization, and uses the index to change adaptively the mesh. We have shown AMR to be very effective in two and three dimensions, often gaining a factor of 10 to 100 in computational efficiency. Work has begun on parallel implementations of AMR. A second goal is to combine AMR with Composite Overlapping Grid (COG) techniques.

We are also interested in the reconstruction of signals from data. One focus is applications to tomography—that is, the nondestructive determination of the internal structure of a physical object. Recent work includes the derivation of a new method for calculating the mass in a tomographic slice; development of a highly efficient ray tracing code to provide reconstruction data sets; and derivation of a method to use the predicted spectrum in designing detectors optimized for the true spectrum.

A related research topic is the use of wavelet transforms. Wavelets are particularly appropriate for applications involving nonstationary processes or highly irregular or nonperiodic data sets. Current work aims at developing novel wavelet techniques for solving PDEs, and data compression methods for output of large-scale simulation models.

Oak Ridge National Laboratory

COMPUTATIONAL SCIENCE EDUCATION PROJECT

Work in this project is toward the development of a syllabus and accompanying teaching materials for early graduate level instruction in computational science. The goal is to establish an interdisciplinary curriculum to bridge the gap between computer science and applied mathematics courses and their applications in the natural and applied sciences and in engineering disciplines. During the first year of the project material was collected into a draft of more than 500 pages representing the contributions of 20 authors from the disciplines of acoustics, aerospace engineering, biology, chemistry, chemical engineering, computer science, earth sciences, mechanical engineering, nuclear engineering, numerical mathematics, oceanography and physics.

Four workshops were held during this year. The first two served the purpose of recruiting authors and outlining writing assignments. The third was held to review and discuss the first partial draft, about 285 pages, and to identify additional areas needing attention. The fourth served the purpose of examining the second draft, about 500 pages, and determining the chapters needed for teaching during the fall of 1992.

The draft has been developed and will be distributed in electronic form to permit fast distribution and easy updates and to eliminate the need to restrict the amount of material included in the product. The structure consists of an introduction to high-performance computing, selected computational methods, and a series of case studies covering acoustics, biology, chemistry, earth sciences, mechanical engineering, nuclear engineering, and physics. Additional material has been discussed with a number of academicians for extensions in the areas of computational fluid dynamics, seismology, atmospheric sciences, particle simulation, and health sciences.

MATHEMATICS OF PHYSICAL SYSTEMS

Research in this area concentrates on the development, analysis, and numerical solution of the partial differential equations that model complex phenomena of interest to DOE. Current problems of interest are those associated with the study of environmental remediation and similar applications. The major tasks are the following. (1) The numerical modeling of groundwater flow and transport involves the solution of the fluid flow field in porous media and the advection dominated convection-diffusion transport equations. The associated linear equations are frequently solved by domain decomposition with preconditioned conjugate gradient acceleration. Current investigations focus on developing a more efficient global "probe" preconditioner. (2) Motivated by a major In-Situ Vitrification (ISV) test conducted at ORNL in 1991, a 3-D simulation code modeling an ISV process has been developed. The model incorporated heat transfer and electric field calculation, together with a means for accounting for the effects of natural convection in the melt. (3) Research has begun on applying optimal control theory to the groundwater effort to produce solutions for controlling realistic large scale problems. Optimal control of bioremediation is also under investigation. (4) The application of continuum-scale models for modeling subsurface flow requires appropriate macroscopic transport parameters characterizing the soil. Fractal structures present in heterogeneous soils should be modeled, and their effect on the large-scale flow parameters studied. For these problems, boundary integral methods can be a highly efficient alternative to finite element algorithms. Preliminary experiments have been encouraging, and further testing of this method for flow and wave scattering problems will be carried out. The effort required to implement the hypersingular equations has been greatly reduced by developing symbolic computation techniques for evaluating the singular integrals.

NUMERICAL ANALYSIS AND SCIENTIFIC COMPUTING

Research in this area focuses on the development and analysis of basic algorithms and techniques for solving large scale scientific problems on advanced computer architectures. The principal areas of research are the direct and iterative solution of large sparse systems on both serial and parallel computers, and parallel algorithms for solving partial differential equations. This project also supports the Householder Postdoctoral Fellowship in Scientific Computing. New and improved algorithms have been designed, developed, and implemented for solving large sparse linear systems by direct methods. Our effort on iterative methods for sparse linear systems has resulted in important bounds on the decay entries in incomplete factorization and a simplified, less expensive ordering method has proven effective for large computational stencils arising in solving Navier-Stokes equations. A rearrangement of the computation in the conjugate gradient method has been discovered that allows the computation of two separate inner products of distributed vectors at once, requiring only one communication phase. A theoretical correspondence has been established between epsilon-pseudo-eigenvalues and Fourier analysis applied to the iterative solution of nonsymmetric matrices, partially explaining the utility of the Fourier analysis. Algorithms for common parabolic and hyperbolic PDEs have been developed that can be parallelized in time, achieving a polylogarithmic parallel complexity, while retaining linear serial complexity. The new algorithm for parabolic problems has promise as a practical algorithm for massively parallel multiprocessors.

DATA HANDLING AND ANALYSIS

Modern approaches to data collection and generation lead to complex and often massive distributed/parallel databases from which information and scientific insight must be drawn. The ARM database to be located at ORNL is one example of such a database. Other examples include large computer-generated data sets generated by Grand Challenge activities. Effort in this project melds our pervious work in the design and analysis of computer experiments with the development of new visualization techniques necessary for exploring such databases. Techniques for designing and analyzing computational experiments applicable to inverse problems that arise in Grand Challenge problems are addressed. Tools for use with models described by time-dependent partial differential equations such as visualization, scale-up, and parameter estimation are developed and applied to important models of fluid flow problems and *in situ* vitrification processes.

Sandia National Laboratories

APPLIED MATHEMATICS, COMPUTING SCIENCE, AND EDUCATION

Sandia's continuing research program in BRHR has components in applied/computational mathematics, computing science, and education. Our applied mathematics research focuses on the development of scalable algorithms for high-performance computers: software and algorithms are as important to the speed and utility of high speed computing as the improvements in the hardware itself. Our computing science research focuses on the development of the software tools and related technology needed to enable MP high performance computing systems to achieve wider acceptance and everyday use by government and industry. Our educational activities included the Adventures in Supercomputing Program (reported in Section 2) and a Research Fellowship Program.

MP Algorithms for Linear Systems

The solution of linear systems underlies most problems in scientific computation. Sandia's program encompasses scalable methods for both dense and sparse systems and includes both direct and iterative methods. For dense systems, Sandia has developed a suite of high-performance direct solvers, including LU and QR factorizations; these routines have achieved scaled speedups of nearly 1000 on a 1024-processor nCUBE2. Current research focuses on out-of-core methods and methods for MP condition number estimation. For large sparse systems, highly-parallel multifrontal solvers for applications arising from finite-difference or finite-element models are under development. A Sandia developed suite of preconditioned, Krylov-based iterative solvers for symmetric and nonsymmetric systems achieved scaled speedups of over 900 and sustained computation rates in excess of 1 Gflop on a 1,024-processor nCUBE2 for a wide variety of applications. Research continues on Krylov-based methods, focusing on the quality of the preconditioner and the speed of the sparse matrix-vector multiplication.

MP Methods for Nonlinear Systems and Optimization

Most important problems in scientific computation are nonlinear, including the Navier-Stokes equations for fluid flow, most optimization problems, and many inverse problems. Nonlinear problems are generally computationally intensive and hence are natural candidates for MP computing. The initial focus for Sandia's research in this area has been on techniques for solving a prototypical inverse problem (inverse optical scattering) and a molecular conformation problem (analysis of molecules for cancer drugs).

Load Balance and Mapping Methods

An important and pervasive problem in MP computing is the assignment of parts of a large problem to processors to achieve optimal parallel efficiency. Efficient load balance algorithms for finite difference, finite element, particle-in-cell, and other types of scientific computations can have a significant impact on complex applications solved on MP computers. Sandia researchers have generalized and improved an existing spectral method to allow unbalanced computation and communication loads, to minimize not only the amount of communication between processors but also the number of messages traveling between distant processors in a hypercube or mesh architecture, and to partition the problem into eight pieces at each stage of recursion instead of two. The modified method has been combined with the well-known Kernighan-Lin heuristic to produce an efficient, reliable, and robust algorithm that has been tested extensively on a wide variety of applications. An MP version is currently under development.

MP Adaptive Refinement Methods for Finite Element and Finite Difference Calculations

The development of highly parallel adaptive methods for large-scale finite element and infinite difference calculations can play a significant role in addressing the Grand Challenges. Sandia has developed a fast, accurate, and robust serial implementation of adaptive refinement methods in 2-D that runs on both the nCUBE2 and the Intel iPSC860 systems. The resulting software package has been successfully implemented in a new application code at Sandia for modeling ignition in multi-phase granular materials.

Sandia is also developing high-order, finite-element based, adaptive procedures for the solution of hyperbolic conservation laws. These methods will more easily handle problems with complicated geometries and boundary conditions and be implemented more efficiently on MP computers than traditional high-order finite difference schemes. Experiments with 1-D and 2-D problems have demonstrated the method's high-order accuracy in smooth solution regions and the sharp capture of discontinuities. Parallel efficiencies of over 98% have been achieved on the nCUBE2. The method is being extended to 3-D and dynamic load balancing procedures will be incorporated .

Nonlinear Problem in Combustion Theory

Our combustion research is centered on the mathematical modeling of significant nonlinear phenomena arising in the study of flame propagation in gases, solids, and supercritical fluids; in the deflagration of multiphase systems such as solid and liquid propellants; and, most recently, in nonlinear acoustic phenomena arising in practical combustion devices such as pulse combustors used for heating and drying applications. Our goal is to predict and understand bifurcation, or transition, phenomena that are common to such problems. Techniques in nonlinear stability and bifurcation theory are employed to achieve a reduction of the governing system of partial differential equations to dynamical systems of evolution equations for the amplitudes of appropriate eigen-modes. Our analyses of such evolution equations has been highly successful in predicting and explaining new routes to quasi-periodicity and chaos stemming from both resonant and nonresonant inode interactions.

Data Visualization for MP Supercomputer Applications

Visualization is critically important because computations on MP computers today typically generate many Gigabytes of data. Sandia researchers have developed an MP marching cubes algorithm for generating polygonal isosurfaces directly on distributed memory MP machines. Previously, isosurfaces were generated on a graphics workstation, limiting the size of datasets to the memory size of the graphics workstation. Currently, the isosurfaces are transferred to a workstation for rendering, thereby graphically displaying the results of simulations that fill the Gigabytes of memory in MP supercomputers. One algorithm was recently integrated with both MP shock physics and fluid dynamics application codes to create high fidelity movies of the simulations. An environment was built to demonstrate these MP capabilities. Sandia researchers recently demonstrated the ability to drive a frame buffer remotely at HiPPI speeds over a 19-kilometer fiber-optic link. Current research will provide users of parallel computers with a visualization environment that supports large data sets, and both real-time and post-processing of data.

Automatic Mesh Generation

In many important computations (structural dynamics, aerodynamics, hydrodynamics), the time required to generate a computational mesh is a significant fraction of the total solution time. To reduce this time, Sandia is developing tools that generate 2-D and 3-D meshes of much higher quality in much less time than currently possible; the research will also provide adaptivity tools to allow more accurate and efficient analysis of mechanics problems. A PAVING algorithm is under development that will mesh a 2-D or 3-D surface with an all-quadrilateral or hexahedral mesh with only the geometry of the surface and the approximate size of the elements as input. A PLASTERING algorithm is also under development that will fill a 3-D solid with an all hexahedron mesh with the geometry being specified by using a solid modeling software package. While this is even more challenging than the PAVING algorithm, results to date for PLASTERING have been extremely promising.

Operating System Research and Related Issues

A team of University of New Mexico (UNM) and Sandia researchers has recently developed a portable, UNIX-based operating system especially tailored for distributed memory MP supercomputers. The operating system is part of a research program to test novel ideas in scalable MP systems software: high-speed asynchronous interprocessor message passing, automatic load balancing, improved performance of large-scale applications, and interoperation among heterogeneous MP machines. The research is directed by Sandia staff and involves 5–10 UNM graduate students each quarter. To date, significant advances in asynchronous message-passing protocols permit run-time construction of arbitrary message patterns: once constructed, the message patterns are repeatedly invoked with minimal cost, thereby, improving performance of many iterative application algorithms. The operating system kernel itself is novel because of its very small size and high efficiency. Two US companies are interested in including the system kernel in their future MP products.

Performance Visualization and Debugging for MP Supercomputers

Graphical performance monitoring tools are crucial to analyzing processor activity (percentage of time spent in computation) and interprocessor communications when using hundreds and thousands of processors. Sandia has developed the first such tool that provides a real-time display. The tool is portable and has a minimal impact on application performance. Researchers are using the performance tool to identify hot spots in computations, improve load balance, and minimize communication time. The research focuses on the person-machine interface. New 2-D and 3-D performance displays are created and tested by application code developers. In addition, the performance monitor is being integrated with an MP debugger to help reduce the time required to write and debug MP applications software.

Applied Mathematical Sciences Research Fellowships

AMS Fellowships at Sandia provide an excellent opportunity for innovative research in scientific computing on advanced architectures, and promote the transfer of technology from the laboratory research environment to industry and academia through the advanced training of computational scientists. Candidates must be US citizens, must have earned a recent Ph.D. degree or the equivalent, and must have a strong interesting advanced computing research. The fellowship appointment is for a period of one year, and may be renewed for a second year. The fellowship program was instituted in FY89.

University Summaries

Mathematics of Physical Systems

American Mathematical Society

Wavelets and Applications
James Maxwell

This grant provides travel and subsistence support for participants in a Summer Research Conference on Wavelets that is sponsored by the American Mathematical Society, the Institute of Mathematical Statistics, and the Society for Industrial and Applied Mathematics. The conference, which is devoted to one of the most exciting new areas of applied mathematics, numerical analysis, and signal and image processing, will bring together theoreticians, applied mathematicians, applications' people and people from industry to learn about the latest results in wavelet theory and applications of wavelet techniques.

American Mathematical Society

Exploiting Symmetry in Applied and Numerical Analysis
James Maxwell and Kurt Georg

This grant provides participant support costs attendant to a Summer Seminar Series on Symmetry that is sponsored by the American Mathematical Society and the Society for Industrial and Applied Mathematics. The goal of the seminar is to bring together applied mathematicians and numerical analysts in an effort to expose them to the wealth of ideas in the areas of bifurcation theory, stability theory, and continuation theory for partial differential equations that exploit the inherent symmetry of physical and biological systems.

University of Arizona

Singularities and Symmetries of Nonlinear Ordinary and Partial Differential Equations

Michael Tabor

The principal investigator and his colleagues are studying the singularity structure and the symmetries of solutions of nonlinear ordinary and partial differential equations. For ordinary differential equations the researchers will analyze various tests for integrability that are based on the work of Painleve, whereas for partial differential equations they will attempt to develop similar tests by using ideas from the theory of several complex variables. The insights into the behavior of solutions of nonintegrable systems gained from these studies will be applied to understanding the nature of chaotic and turbulent motions.
Students: 1 postdoc, 2 graduate.

California Institute of Technology

Analysis, Scientific Computing and Fundamental Studies in Fluid Mechanics

Herbert Keller and Philip Saffman

Daniel Meiron, co-investigator

The principal investigators and their colleagues are receiving continuing support for their project in applied mathematics/scientific computing that combines the development and implementation of numerical methods for the simulation of fluid dynamical phenomena with the development of software for general scientific computing and the analysis of basic fluid dynamics problems related to vortex phenomena in viscous and inviscid fluids. In many respects this work represents an ideal project for sponsorship by the base program in applied mathematics because it features a genuine interaction among the core activities of mathematical modeling, theoretical analysis, and numerical simulation of important physical phenomena on parallel architectures.
Students: 1 postdoc, 4 graduate.

University of Chicago

Studies of Complexity in Fluid Systems

Leo Kadanoff and Peter Constantin

The principal investigators and their collaborators from the mathematics and physics departments will study a number of interesting problems in fluid dynamics that involve interfacial phenomena, the behavior of preturbulent and turbulent flows, and the evolution of coherent fluid structures using analytical, numerical,and experimental techniques. These various phenomena exhibit what is termed "self-organized critical behavior." In this context "critical" signifies that events of all sizes occur in the system, with different physical laws operating on different length or time scales, while "self-organized" implies the existence within the system of feedback mechanisms that produce marginally-stable steady-states. The researchers on this project will focus on self-organized critical behavior as it arises in the transition from laminar to turbulent flow and in the evolution of fully-developed turbulence.
Students: 2 postdoc, 2 graduate.

Harvard University

Mathematics and Physics

Arthur Jaffe and S.-T. Yau

The proposed research continues work on the interaction of modern mathematics and physics in the general areas of quantum field theory, string theory and the geometry of space-time. In particular, the principal investigators are studying various fundamental questions in the geometry of infinite-dimensional spaces by means of noncommutative differential geometry, the geometry and topology of quantum groups, vacua in superstrings and their geometry, the solution space of the Yang-Mills equations, topological quantum field theory, and nonperturbative string theory.
Students: 3 postdoc, 2 graduate.

BEST COPY AVAILABLE

University of Houston

Shock Stability in Systems that Change Type

Barbara Keyfitz

In this project the principal investigator is studying various questions related to the stability of solutions of nonstandard conservation laws of the type that describe the behavior of multiphase flows and flows in porous media. In particular, she is investigating the properties of shock wave solutions that satisfy generalized entropy conditions and the symmetry structure of solutions of conservation laws in higher dimensions. This work provides the theoretical underpinning for DOE-supported modeling and computational projects concerned with groundwater flow and enhanced oil recovery.
Students: 1 graduate.

Indiana University

Approximation of Attractors and Applications

Roger Temam and Ciprian Foias

The principal investigators and their colleagues are studying the mathematical structure of turbulent and chaotic flows governed by the Navier-Stokes equations. Turbulent flows are prime examples of dynamical systems with infinitely many degrees of freedom, but many of their properties can be understood by analyzing certain finite-dimensional objects that in some sense characterize the infinite-dimensional turbulent structures. The principal investigators are two of the pioneers in taking this approach to studying the long-time behavior of solutions of the Navier-Stokes equations and developing accurate numerical methods for approximating the finite-dimensional structures in turbulent flows.
Students: 1 postdoc, 1 graduate.

Louisiana State University

Conference on Evolution Equations

James Dorroh and Jerome Goldstein

This grant supports a conference on evolution equations that is bringing together researchers and graduate students who are using the tools of modern functional analysis to solve interesting and physically important classes of partial differential equations. While the emphasis of the conference is on presenting and discussing the latest results in the theory of abstract partial differential equations, the conference itself serves to attract mathematicians and physical scientists with a common interest in applying analytical methods to solve problems in mathematical physics of direct relevance to the applied mathematics program at DOE.

University of Maryland at Baltimore County

Hyperbolic Heat Transfer Problems with Phase Changes and Nonlinear Transport

James Greenberg

This research focuses on modeling the behavior of solutions of conservation laws of the type that arise in various continuum and transport processes by means of discrete velocity models of the Boltzmann equation. Recent results include one that settles a long-standing conjecture of Peter Lax regarding the appropriateness of using a discrete particle model to approximate the behavior of continuous flows with shock wave discontinuities.

Massachusetts Institute of Technology

Research in Geometry and Elementary Particle Physics
Isadore Singer

Scott Axelrod, co-investigator

The principal investigator and his colleagues are working on the interface between modern mathematics and modern physics. Among the topics currently under investigation are the analysis of various models in quantum field theory using geometric perturbation theory; the relationship between methods based on perturbation theory and methods based on combinatorial topology for studying knot invariants; and a number of problems in current algebras, quantum gravity, and quantum groups. This project is co-funded at DOE by the High-Energy Physics Program.
Students: 3 postdoc, 3 graduate.

University of Minnesota

Environmental Studies: Mathematical, Computational and Statistical Analysis
Avner Friedman and Willard Miller

This grant provides partial support for a summer workshop at the Institute for Mathematics and Applications that brings together researchers from mathematics, computer and computational science, statistics, chemistry, physics, geology, meteorology and environmental science to review ongoing work and preview future directions in the complementary areas of environmental modeling and simulation, environmental data assimilation, stochastic modeling and optimization, and global climate modeling. The grant is co-funded at DOE by the Office of Health and Environmental Research.

National Academy of Sciences

Core Support for the Board on Mathematical Sciences
John Lavery

The Department of Energy contributes partial support for the core activities of the Board on Mathematical Sciences. These core activities include the work of the Board in organizing the annual Department Chairs Colloquium and the National Science and Technology Week Symposium; in preparing reports throughout the year on topics in research, education and cross-disciplinary science of importance to the mathematical sciences community; and in advising the federal agencies on new directions in mathematics research, education and outreach that originate from the community or from other disciplines. Contributions to this effort are made also by the other ICEMAP agencies (AFOSR, ARO, DARPA, NSA, NSF and ONR).

Northwestern University

Studies in Nonlinear Problems of Energy
Bernard Matkowsky

Alvin Bayliss and Zeev Schuss, co-investigators

The principal investigator and his colleagues are studying problems in bifurcation theory and nonlinear stability theory that are important in various energy-related applications using a combination of analytical and numerical techniques. In particular, some of the current problems involve the transition from laminar to turbulent flame propagation in gaseous combustion, gasless solid fuel combustion, and supercritical combustion.
Students: 1 postdoc.

University of Washington

Nonlinear Resonance
Jirair Kevorkian

The principal investigator continues his work on the analysis and interpretation of resonant behavior in dynamical systems by means of various asymptotic techniques. In particular, he will study the behavior of slowly varying nonlinear oscillatory systems, with special emphasis on systems that model the operation of free-electron lasers and systems that exhibit "bursting" phenomena typical of nerve cells.
Students: 1 graduate.

Yale University

Research in Geometry, Symmetry, and Physics
Igor Frenkel, Howard Garland, and Gregg Zuckerman

Feza Gursey, Gregory Moore, and Hubert Saleur, co-investigators

This grant provides postdoctoral, graduate student, and visiting scientist support for work that lies at the interface between mathematics and physics. Specifically, this group of mathematicians and physicists will study how the latest advances in geometry, topology, and representation theory can further our understanding of basic phenomena in quantum field theory, quantum gravity, and string theory. Alternatively, it is clear that investigations into problems involving low-dimensional field theories are linked closely with recent work in algebraic geometry, Lie and von Neumann algebras, knot theory, and the geometric theory of nonlinear partial differential equations.
Students: 1 postdoc, 2 graduate.

Optimization Theory and Mathematical Programming

Columbia University

Algorithms for Mathematical Programming
Donald Goldfarb

The principal investigator will be developing and implementing algorithms for solving some important classes of mathematical programming problems; in particular, he will consider simplex algorithms and interior-point algorithms. The work with simplex algorithms is directed at solving network flow problems like maximum flow problems, minimum cost flow problems, and multicommodity flow problems. The work with interior-point algorithms is directed at solving specially-structured linear programming problems, studying large-step path following methods, and extending the linear theory to nonlinear programming problems. Finally, the principal investigator will combine these two approaches in attacking various classes of difficult nonlinear programming problems.
Students: 1 graduate.

Cornell University

New Algorithms for Large-Scale Linearly-Constrained Optimization with Bound Constraints
Thomas Coleman

The principal investigator will develop and implement passive-set algorithms for solving large minimization problems subject to linear constraints on parallel architectures. Passive-set algorithms are found to possess several properties that make them suitable for parallel implementation; namely, they tend to converge with very few iterations, parallelism and sparsity can be exploited in a black-box mode, and any improvements in large-scale linear algebra solvers can be used to speed up these algorithms with little or no reprogramming.
Students: 1 graduate.

University of Maryland at Baltimore County

Research on Primal-Dual Interior Point Algorithms

Yin Zhang

This grant funds a graduate student to work with the principal investigator on the primal-dual interior point method for solving large-scale linear programming problems. In particular, the researchers will combine analytical and numerical methods to develop and implement a robust code on high performance platforms to study various large programming problems from the theory of communications networks and resource allocation models.
Students: 1 graduate.

Northwestern University

Optimization and Eigenvalue Computations with Application to Meteorology and Oceanography

Jorge Nocedal and Valerie Taylor

The principal investigators will work on developing and implementing faster and more robust numerical algorithms for the solution of large-scale nonlinear optimization and eigenvalue problems on massively parallel architectures. The particular nonlinear problems posed for study arise in analyzing various data assimilation models for reducing the data sets associated with oceanographic and atmospheric measurements and in minimizing various types of nonlinear functionals.
Students: 2 graduate.

Rice University

Research on Trust-Region Algorithms for Nonlinear Programming

Richard Tapia and John Dennis

The principal investigators are developing, implementing, and testing robust global algorithms for the efficient solution of general linear and nonlinear programming problems. In particular, they are using interior-point methods to handle inequality constraints in programming problems and trust-region algorithms to attack unconstrained problems, as well as the large sparse problems that arise in oil-reservoir simulations, and they are formulating a unified approach to the convergence theory of trust-region methods applied to nonsmooth optimization problems. Recently this group has implemented a novel secant method for constrained optimization problems that outperforms existing secant algorithms. The principal investigators are also actively involved in recruiting and mentoring minority and women graduate students and integrating them into their research program.
Students: 2 graduate.

Stanford University

Solving Linear Problems Under Uncertainty: Using Decomposition, Importance Sampling and Parallel Processors

George Dantzig and Peter Glynn

Gerd Infanger, co-investigator

The principal investigators and their colleagues are continuing their analytical and numerical work on stochastic linear programming problems, which are linear programming problems in the presence of uncertainty. Stochastic linear programming is a difficult area of theoretical and practical importance. The progress made in this project will have implications for the efficient generation and distribution of electric power over a large grid, the efficient scheduling and transmission of telephone calls over a large network, and the determination of optimum strategies for airline crew and flight scheduling. Apart from its practical significance, this work in stochastic linear programming has relevance for addressing broader questions related to applying essentially deterministic techniques to resolve nondeterministic phenomena.
Students: 1 graduate.

Stanford University

New Approaches to Linear and Nonlinear Programming
Walter Murray and Michael Saunders

Philip Gill, co-investigator

The investigators on this project will develop and implement fast and efficient algorithms for the solution of large-scale linear and nonlinear programming problems. Such programming problems arise in a host of DOE-supported activities. In this work the investigators will focus on developing robust algorithms for implementation on multiprocessor and massively parallel architectures to solve large power generation models that accurately describe actual energy production systems in operation today.
Students: 1 graduate.

Dynamical Systems Theory and Chaos

University of California at Berkeley

Stability, Bifurcation and Control of Hamiltonian Systems
Jerrold Marsden

Tudor Ratiu, co-investigator

The principal investigator and his colleague at UC Santa Cruz are working in the area of nonlinear stability, bifurcation, and control of mechanical systems with symmetry. Theirs is a broad, ongoing research effort that involves the application of powerful tools from a number of different areas of mathematics to the study of conservative dynamical systems with finitely and infinitely many degrees of freedom. Particular topics proposed for study include the separation of rotational and internal modes in rotating Hamiltonian systems, symmetry-breaking bifurcations in Hamiltonian systems, variational principles for the Hamiltonian-Jacobi equations, and the application of ideas from gauge theory to optimal control problems for symmetric mechanical systems.
Students: 1 graduate.

University of Maryland

The Study of Effects of Small Perturbations on Chaotic Systems
Celso Grebogi and James Yorke

Edward Ott, co-investigator

This ongoing research project is focusing on the chaotic dynamics of dissipative nonlinear systems in the following two main areas: the control of chaotic processes using periodic orbits and perturbation methods that take the state of the chaotic system to a specified target region in phase space and the creation of typical trajectories of chaotic systems in which the statistics of the system are accurate over a long time interval. Results in these areas have increased our understanding of the dynamic responses of systems that describe physical phenomena like the operation of chemical reactors, and this understanding has led to the development of practical methods for controlling chaotic behavior in real time.
Students: 1 postdoc, 2 graduate.

Transient chaos in the Lorenz system.

Geometric and Symbolic Computation

Florida Institute of Technology

The Use of Symbolic Computation in Radiative, Energy and Neutron Transport Calculations
Jay Frankel

The principal investigator will study solutions of integral and integro-differential equations that describe radiation and neutron transport by means of numerical techniques based upon the use of symbolic computation. Symbolic computation will be employed to perform analytical studies of the solution sets of these equations and to aid in the numerical approximation of solutions in cases where straightforward numerical approaches run into difficulties. The problems to be addressed arise in a number of DOE-supported activities at the national laboratories and at universities. The research to be conducted will improve our understanding of the basic properties of solutions of integral equations, and it will lay the groundwork for developing efficient numerical algorithms for solving real problems in neutron transport and plasma physics.

University of Massachusetts at Amherst

Computation and Graphics in Mathematical Research
David Hoffman and Joel Spruck

William Meeks, David Oliver, and Robert Kusner, co-investigators

This project involves the integration of pure and applied mathematics, high performance computers, and advanced interactive computer graphics into an environment where researchers can explore important scientific problems using the tools of high-speed computation and visualization. Among the many problems currently under investigation are the global geometric properties of minimal surfaces, free boundary problems from plasma physics and fluid dynamics, the geometry of polymers and genes, and the use of geometric tools to manage large data sets.
Students: 1 graduate.

University of Minnesota

Center for Computation and Visualization of Geometric Structures
Albert Marden

This grant helps to support the activities of the Center for Computation and Visualization of Geometric Structures at the University of Minnesota, which is one of the Science and Technology Centers funded by the National Science Foundation. The Center has more than exceeded everyone's expectations in providing a home for outstanding research and teaching in the areas of geometric computation and scientific visualization. In addition to fostering the research of investigators on campus, the Center serves as a focal point for research conferences and workshops, educational and outreach activities, and as an ideal location for attracting long-and short-term visiting mathematicians, engineers, and scientists from around the world who are the leaders in the growing field of geometric computation.
Students: 2 postdoc.

Numerical Analysis and Scientific Computation

University of California at Los Angeles

Center for Parallel PDE Algorithms
Tony Chan

Tarek Mathew and Beth Ong, co-investigators

The Center for Parallel PDE Algorithms is an umbrella organization on the UCLA campus comprised of a group of faculty doing research in computational science on a variety of small parallel computers located at UCLA. The goal of the research project is to develop algorithms based on the concept of domain decomposition and software tools based on iterative methods for the fast and efficient solution of large-scale scientific problems on multiprocessor–processor and massively parallel architectures. Domain decomposition techniques lend themselves naturally to implementation on parallel machines, and they generate powerful algorithms that are able to resolve phenomena featuring behavior on disparate spatial scales. The computation of large-scale nonlinear problems creates a need for fast and efficient

linear algebra solvers, which the investigators are addressing by developing and implementing robust parallel iterative methods.
Students: 1 postdoc, 1 graduate.

University of Colorado at Boulder

Numerical Methods for the Unsymmetric Tridiagonal Eigenvalue Problem

Elizabeth Jessup

In this research project the principal investigator will develop and implement on vector and parallel architectures fast and efficient algorithms for solving unsymmetric eigenvalue problems for large systems of linear equations. Although much of the analytical and numerical work on eigenvalue problems for large linear systems is devoted to symmetric systems, eigenvalue problems for unsymmetric systems arise in a number of important applications of direct relevance to many DOE-supported activities in the base program and in the Grand Challenge work at laboratories and universities. The basic approach, which has the flavor of "divide and conquer" strategies, is to develop efficient parallel methods to compute both eigenvalues and eigenvectors that avoid the numerical instabilities associated with previous implementations of existing solvers on parallel machines.
Students: 1 graduate.

University of Colorado at Denver

Fast Algorithms for Transport Models

Thomas Manteuffel

The principal investigator is developing and implementing algorithms based on multigrid methods for solving accurately and efficiently the integro-differential equations that arise in the theory of neutron transport in one- and two-space dimensions. Standard numerical approaches often fail to resolve the phenomena accurately in physical regimes where the equations become stiff; however, the principal investigator's multigrid algorithms are able to avoid these pitfalls. Recently, in collaboration with researchers at the Los Alamos National Laboratory, the principal investigator provided dramatic speedups in the performance of some large transport codes at Los Alamos by incorporating his multigrid algorithms into the codes.
Students: 1 graduate.

Duke University

High Resolution Numerical Methods for Compressible Multiphase Flow in Hierarchical Porous Media

John Trangenstein

Complex multiphase chemistry, convection-dominated three-dimensional flow, and hierarchies of heterogeneities in porous media are essential features in the modeling of modern enhanced oil recovery and contaminant transport. In order to simulate the fluid flows associated with these problems the principal investigator will develop higher-order finite difference methods, in combination with sophisticated phase behavior algorithms and adaptive mesh refinement techniques, for implementation on high-performance computers.
Students: 1 graduate.

University of Illinois at Urbana

Numerical Methods for Molecular Dynamics

Robert Skeel

This grant supports a graduate student to work with the principal investigator on developing multigrid algorithms for the fast and accurate solution of the Poisson-Boltzmann equation in three dimensions on various high performance platforms.
Students: 1 graduate.

University of Minnesota

Numerical Methods for Differential-Algebraic Equations in Real-Time Integration of Mechanical Systems

Linda Petzold

The principal investigator will work on designing and implementing stable and efficient numerical algorithms for solving systems of differential-algebraic equations. In particular, she will develop algorithms and software that simulate the behavior of complicated mechanical systems in real time on massively parallel computers. In order to achieve the goal of real-time simulation, the numerical schemes must be very stable and highly accurate, and they must permit the efficient control of step-sizes and the automatic detection of stiffness.
Students: 1 graduate.

University of Minnesota

Piecewise Parabolic Methods for Parallel Computation with Applications to Unstable Fluid Flow in 2- and 3-D

Paul Woodward

The principal investigator and his colleagues are engaged in algorithm development for computational fluid dynamics with a new focus on massively parallel computation, which includes extending the piecewise parabolic method to parallel platforms and developing new techniques for interactive, real-time fluid-flow visualization. In particular, the researchers are applying their numerical methods to the simulation of homogeneous, compressible turbulence, and they are investigating the application of the piecewise parabolic method to the equations of magneto-hydrodynamics in two and three dimensions and to a two-dimensional laser fusion code.
Students: 1 postdoc, 2 graduate.

University of Minnesota

Computational Crystal Growers Workshop

Albert Marden

This grant helped to support a workshop at the Minnesota Geometry Center that brought together computational scientists, physicists, and materials scientists to review ongoing work and preview future directions in the area of computational methods and models for understanding the mathematics and physics of crystal growth. By all accounts the workshop was a success, and it was featured in an article in Science on the applications of computers in mathematics.

University of New Mexico

Numerical and Asymptotic Studies of Complex Flow Dynamics

Evangelos Coutsias, Thomas Hagstrom, and Jens Lorenz

The research project involves the study of various important questions about fluid flows governed by the Navier-Stokes equations through the combined use of asymptotic and numerical methods. In particular, the principal investigators and their students will examine the behavior of fluid flows with fast waves by including boundary and acoustic effects, the behavior of compressible flows in exterior geometries, and the vortex dynamics of wall-bounded flows in which the vortices interact with the rigid walls in a strongly nonlinear manner. Such fluid phenomena arise in a number of other DOE-supported activities at laboratories and universities, and a better understanding of complex flows would contribute substantially to basic research in fluid dynamics and technological applications of high-speed flows.
Students: 7 graduate.

State University of New York at Stony Brook

The Modeling of Complex Continua: Fundamental Obstacles and Grand Challenges

James Glimm and Brent Lindquist

Qiang Zhang, co-investigator

The principal investigators and their colleagues are continuing to work on the modeling, analysis, and numerical simulation of the behavior of reacting and nonreacting flows of single-phase and multiple-phase fluids of the type that occur in combustion, groundwater transport, and oil storage and recovery. In particular, the research effort revolves around the following five main areas: adaptive numerical methods for tracking fronts and fluid interfaces; computation of three-dimensional fluid flows in the chaotic regime and statistical modeling of chaotic flows in two dimensions; understanding nonlinear waveforms and wave diffraction patterns in real materials through numerical simulation; using subgrid modeling and front-tracking techniques to model complex chemical phenomena in reacting flows; and resolution of multiple length scales in chemically reacting flows.
Students: 2 postdoc, 4 graduate.

New York University

Applied Analysis and Computational Mathematics

Peter Lax and Marsha Berger

The DOE program at the Courant Mathematics and Computing Laboratory (CMCL) specializes in the interplay between mathematical analysis and scientific computing. We have many ongoing collaborations with other DOE laboratories (Los Alamos and Livermore) and other institutions (IBM, NIST).

One specific research goal is an understanding of complex fluid flows. We are studying the formation, propagation and diffraction of shock waves and have developed sophisticated computational tools for their simulation. We are also studying oscillatory solutions of nonlinear difference equations as a possible analogy to turbulence. An overlapping goal is the prediction of flows in complex geometries. This includes our computational work on the accurate solution of potential and heat flow problems on domains with hundreds of irregularly shaped holes and our work developing nonbody-fitted grid methods for compressible flows.

During the past year we developed a new Hamiltonian formulation of the Euler equations using the so-called magnetization variables. This opens the door for new theoretical approaches to modeling turbulent flows and also leads to improved vortex-type methods for incompressible flow. We completed the first implementation of a structured adaptive 2D mesh refinement algorithm on a massively parallel machine (the Connection Machine). We developed methods for computing bulk properties of complicated composite materials, such as effective thermal conductivity, dielectric constant or electrical conductivity. In collaboration with researchers at Los Alamos and Livermore, the adaptive mesh refinement algorithm for hyperbolic conservation laws has been extended to three dimensions.

Personnel: 5 full-time faculty, a senior research professor, 2 postdocs, and 8 students.

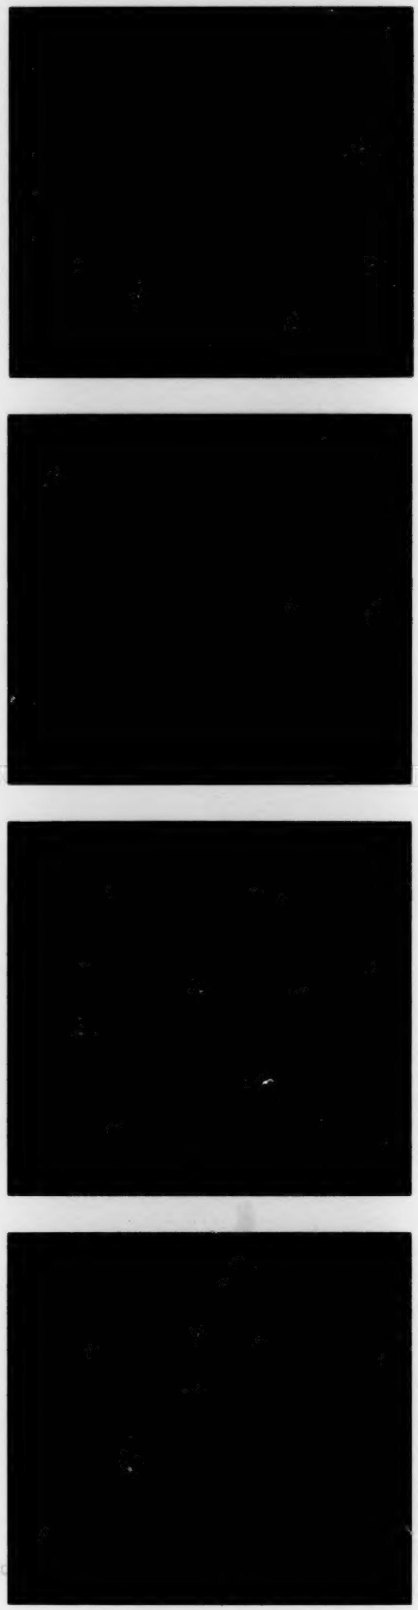

Volume renderings of density during the evolution of a bubble after being hit with a shock wave.

New York University

Domain Decomposition and Parallel Computing for Elliptic Problems
Olof Widlund

The principal investigator is developing and implementing algorithms based on the concept of domain decomposition for solving several classes of elliptic partial differential equations on various parallel platforms. Domain decomposition is an iterative procedure for solving partial differential equations over complicated regions that lends itself naturally to parallel implementation. Specific applications to problems in oil-reservoir modeling, groundwater transport, and structural elasticity are currently under study. The accurate resolution of these phenomena requires the power of massively parallel computers and equally the right software to take full advantage of this parallelism.
Students: 1 graduate.

Pennsylvania State University

Sparse Matrix Algorithms on Distributed Memory Multiprocessors
Alex Pothen

The goal of this research project is to develop algorithms and software for the fast and efficient solution of large sparse systems of linear equations; large least-squares problems; and large optimization problems on parallel, distributed-memory architectures. In particular, the principal investigator will use a spectral nested-dissection algorithm to factor sparse matrices rapidly on parallel machines, and he will design a "proportional mapping" algorithm to assign subtasks to processors in an efficient manner.
Students: 1 graduate.

Rice University

Implicitly Restarted Arnoldi Techniques for Large-Scale Eigenproblems and Linear Systems
Danny Sorensen

The principal investigator is developing and implementing software for the solution of very large nonsymmetric and generalized eigenvalue problems on massively parallel machines. In particular, he will investigate algorithms based on the Arnoldi method, which is a technique for approximating a few eigenvalues of a general matrix at a time, and he will seek to implement the algorithms in a way that avoids the difficulties associated with the large storage demands of classical Arnoldi methods.
Students: 2 graduate.

University of Texas

Parallel Supercomputing: Advanced Methods, Algorithms and Software for Large-Scale Problems
Graham Carey and David Young

David Kincaid and Kamy Sepehrnoori, co-investigators

This project involves the analysis, development and implementation of finite element methods and iterative solution schemes for solving nonlinear partial differential equations on advanced parallel architectures. The research is interdisciplinary in nature, since it involves the Computational Fluid Dynamics Laboratory in the College of Engineering and the Center for Numerical Analysis in the College of Natural Sciences. The present work embraces new developments in finite element techniques for solving partial differential equations of the type that arise in nonlinear mechanics and in iterative methods for solving the resulting systems of algebraic equations, with the emphasis placed on implementing both symmetric and nonsymmetric schemes on state-of-the-art multiprocessor–processor and massively parallel architectures.
Students: 1 postdoc, 3 graduate.

Utah State University

Newton Iterative Methods for Large-Scale Nonlinear Systems

Homer Walker and Kathryn Turner

The principal investigators are using iterative methods based on Newton solvers to approximate the solutions of large systems of nonlinear equations that occur in a number of basic and applied problems in mathematics and engineering. In particular, the researchers will focus on developing and implementing globally convergent iterative methods for systems that arise in discretizing semilinear elliptic partial differential equations.

Wichita State University

Numerical Conformal Mapping: Methods, Applications and Theory

Thomas DeLillo

The principal investigator will develop a number of analytical and numerical methods for improving the applicability of conformal mapping techniques. Given that conformal mapping problems arise in many applied fields such as fluid dynamics and electromagnetics and in the generation of grids in a number of numerical schemes, the proposed research will have an impact on many projects under current support at laboratories and universities. In addition to his theoretical work, the principal investigator will develop conformal mapping software and make it available to the wider scientific community. The software development is driven by studies of fluid flow problems with free boundaries and crack problems in pipes and more general bodies.

University of Illinois

Center for Supercomputing Research and Development

In performance characterization using the Perfect Benchmarks to demonstrate practical parallelism via various efficiency measures, the 32-processor Cedar appears a favorable relative to the 8-processor CRAY Y-MP. A new parallelizing compiler is being built using new techniques for parallel restructuring and tools and methodologies for porting and optimizing programs were developed Development of a second version of the Perfect Benchmarks continued. The Perfect 2 suite will incorporate more contemporary codes with larger datasets drawn from a broader range of application areas.

In cooperation with IMGA-CNR, Modena, Italy, we designed a parallel model of ocean circulation for Cedar. A volume visualization tool has been fitted to the code to provide visualization of the results. The code is presently being used to evaluate the strengths and weaknesses of the Cedar architecture. A longer-term goal is to produce a multimachine code that adapts as a function of architectural parameters and numerical data.

A collaboration with Argonne National Laboratory centered around COMMIX-1AR/P, a three-dimensional transient computer program for thermal hydraulic analysis and WHAMS-3D, a structural dynamics code used for crash-worthiness studies. The COMMIX-1AR/P was able, for the first time, to perform calculations of temperature distribution within a nuclear reactor after a pressurized water reactor steam line break accident, using industrial strength data from Commonwealth Edison.

We made much progress in formalizing a new computational science and engineering program within the College of Engineering at the University of Illinois.

Florida State University

The Supercomputer Computations Research Institute

The Supercomputer Computations Research Institute (SCRI) is a multi-disciplinary research program that has attracted an international group of 40 scientists in disciplines that use high-performance computers. Areas of research include theoretical and experimental high energy physics, nuclear physics, materials science, applied mathematics, computational chemistry, biophysics, meteorology, engineering, geophysics, and economics. SCRI maintains three high-performance computer systems: two of the world's most powerful supercomputers, a Thinking Machines CM-2 Connection Machine and a Cray Y-MP432, and 40+ IBM RISC workstations in a server configuration allow them to work together. These systems and all individual SCRI workstations are interconnected and accessible on a single network.

Major Progress This Year

In experimental high energy physics, SCRI is playing a crucial role in detector development and design for the Superconducting SuperCollider (SSC) presently being built in Texas. Researchers successfully simulated the tracking for the SDC detector, providing valuable information for the choice of the outer tracking system. In theoretical high energy physics, the most recent numerical analysis completed by SCRI's lattice gauge group show that the minimal standard model will describe physics to an accuracy of a few percent up to the energies of the order two to four times the Higgs mass, MH, only if MH 720 50 GeV. If the Higgs has mass as high as 720, then it is expected to be quite broad and be very hard to detect at the SSC. Calculating the low energy properties of QCD continues.

Technology Transfer to Industry

We released two software products during 1992: SciAn and DQS. SciAn is a three-dimensional visualization software package for animating scientific data on computer workstations. We have been using SciAn for modeling and visualizing oil spill simulations, thunderstorm formations, DNA gel electrophoresis, enzyme molecules, and other phenomena. DQS software efficiently and easily "clusters" workstations to create a high-performance computer system as powerful as a supercomputer, but at about one-twentieth the cost.

This year SCRI sponsored a major Workshop on Cluster Computing and a conference on High-Performance Computing and Grand Challenges in Structural Biology to share new ideas with researchers in industry. SCRI is also collaborating with Centel of Florida and IBM on Project Super Star to link SCRI, Centel, Florida A&M University (FAMU), and FSU's Innovation Park with a high-speed fiber optic network.

Educational Outreach

SCRI's scientists and staff are very active in outreach to K-12 education. This past summer we offered once again "The Science Connection From Supercomputer to PC," to introduce the concepts of modeling and simulation using high-performance computing techniques to high school science teachers throughout the state of Florida.